Practical Electronics

The UK's premier electronics, computing and maker magazine

Circuit Surgery
Processing signals digitally

All About Capacitors
Everything you need to know to pick the right capacitor for the job

60W DC load review
A low-cost device for testing supplies & batteries

Raspberry Pi Clock Radio & Media Player

Secure Remote Switch with up to sixteen keyfob remotes

RESET
SD Card Config:
OPTION SDCARD GP13,GP10,GP11,GP12
2N7000
TOSHIBA

The VGA PicoMite: A simple but capable BASIC computer!

PLUS!

Techno Talk – Tomorrow's World, today

Cool Beans – Driving LEDs with serial-to-parallel latches

Net Work – F. J. Camm, the practical man

Dec 2024 £5.99

www.electronpublishing.com | @practicalelec | practicalelectronics

Sysblocks

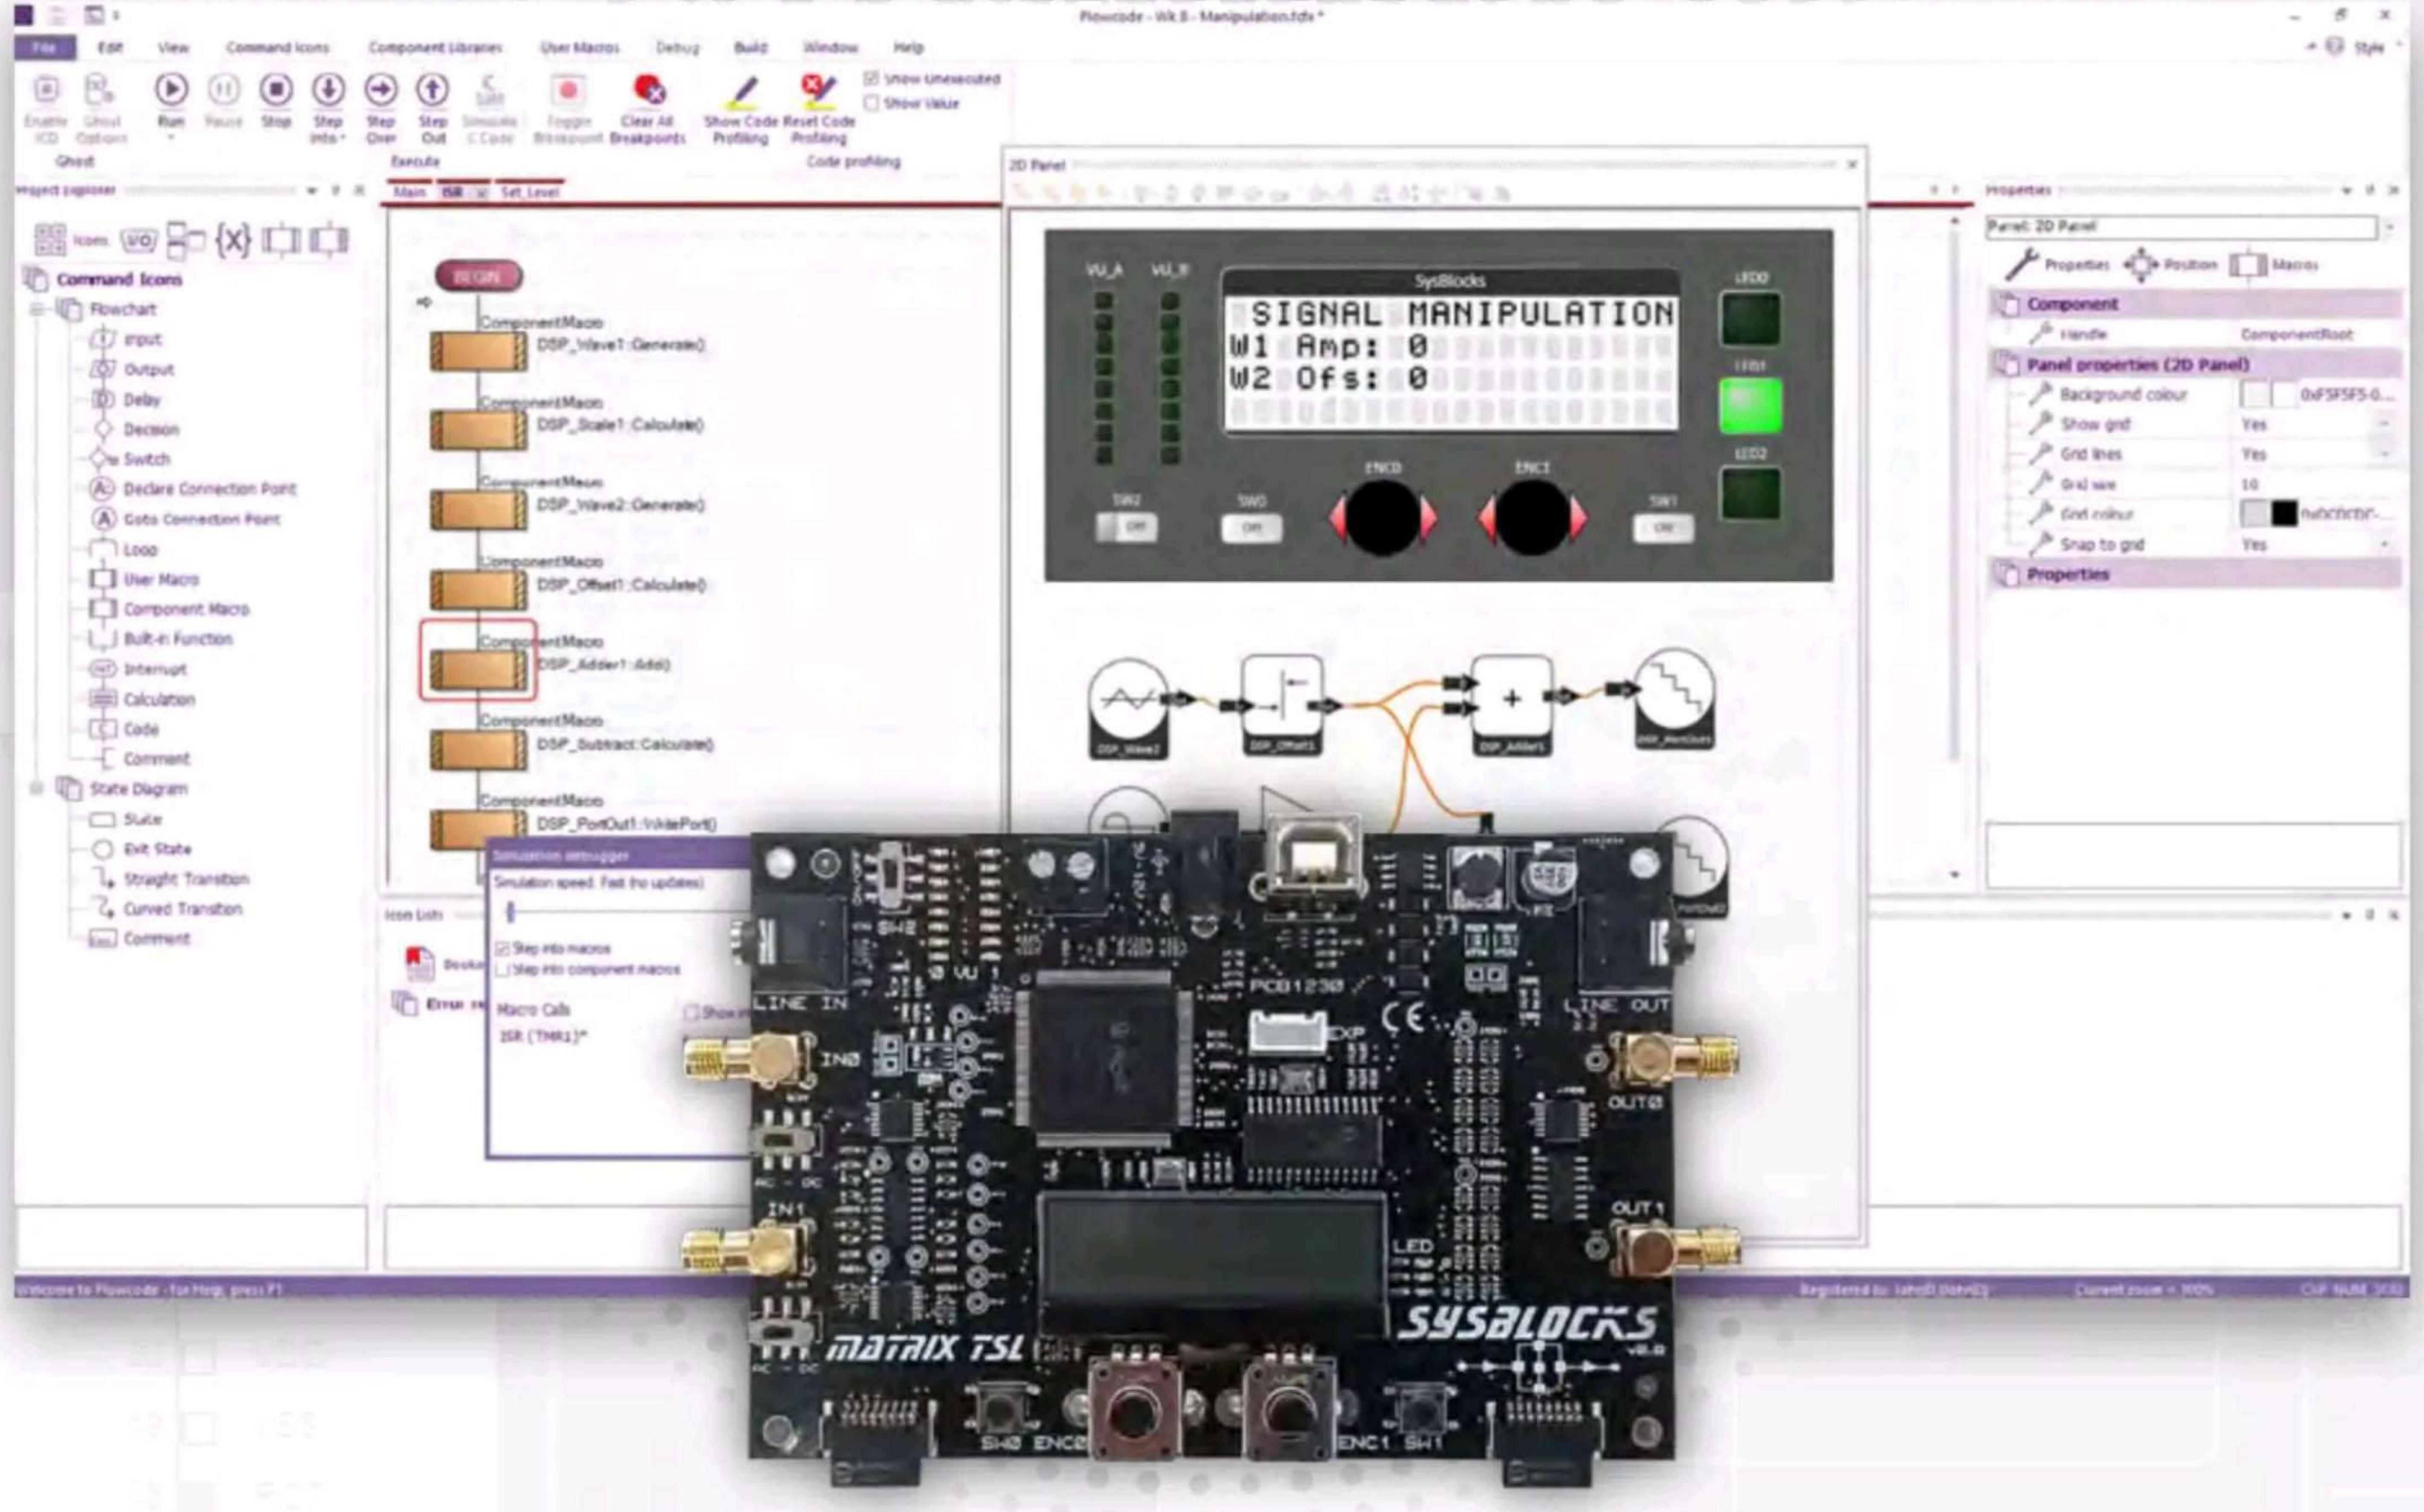

CREATE HIGHLY FUNCTIONAL ELECTRONIC SYSTEMS USING PIC32 TECHNOLOGY TO MANIPULATE ANALOGUE SIGNALS USING FLOWCODE 10.1 AND SYSBLOCKS. FANTASTIC FOR DIGITAL AUDIO APPLICATIONS.

Flowcode Pro, PIC32 pack with Sysblocks board, cables and accessories bundle £610 + VAT
Flowcode with PIC32 pack with Sysblocks board, cables and accessories bundle £330 + VAT

10% off for PE readers using code: **PESYS10**

FLOWCODE 10.1 WITH MORE THAN 20 MICROCONTROLLER TARGETS INCLUDING ARDUINO UNO, PIC18877, RPI2B, AND ESP32 IS NOW FREE FOR HOBBYISTS.

☎ +44 (0) 1422 252 380 — ✉ sales@matrixtsl.com

www.flowcode.co.uk

Practical Electronics

Volume 53. No. 12
December 2024
ISSN 2632 573X

Contents

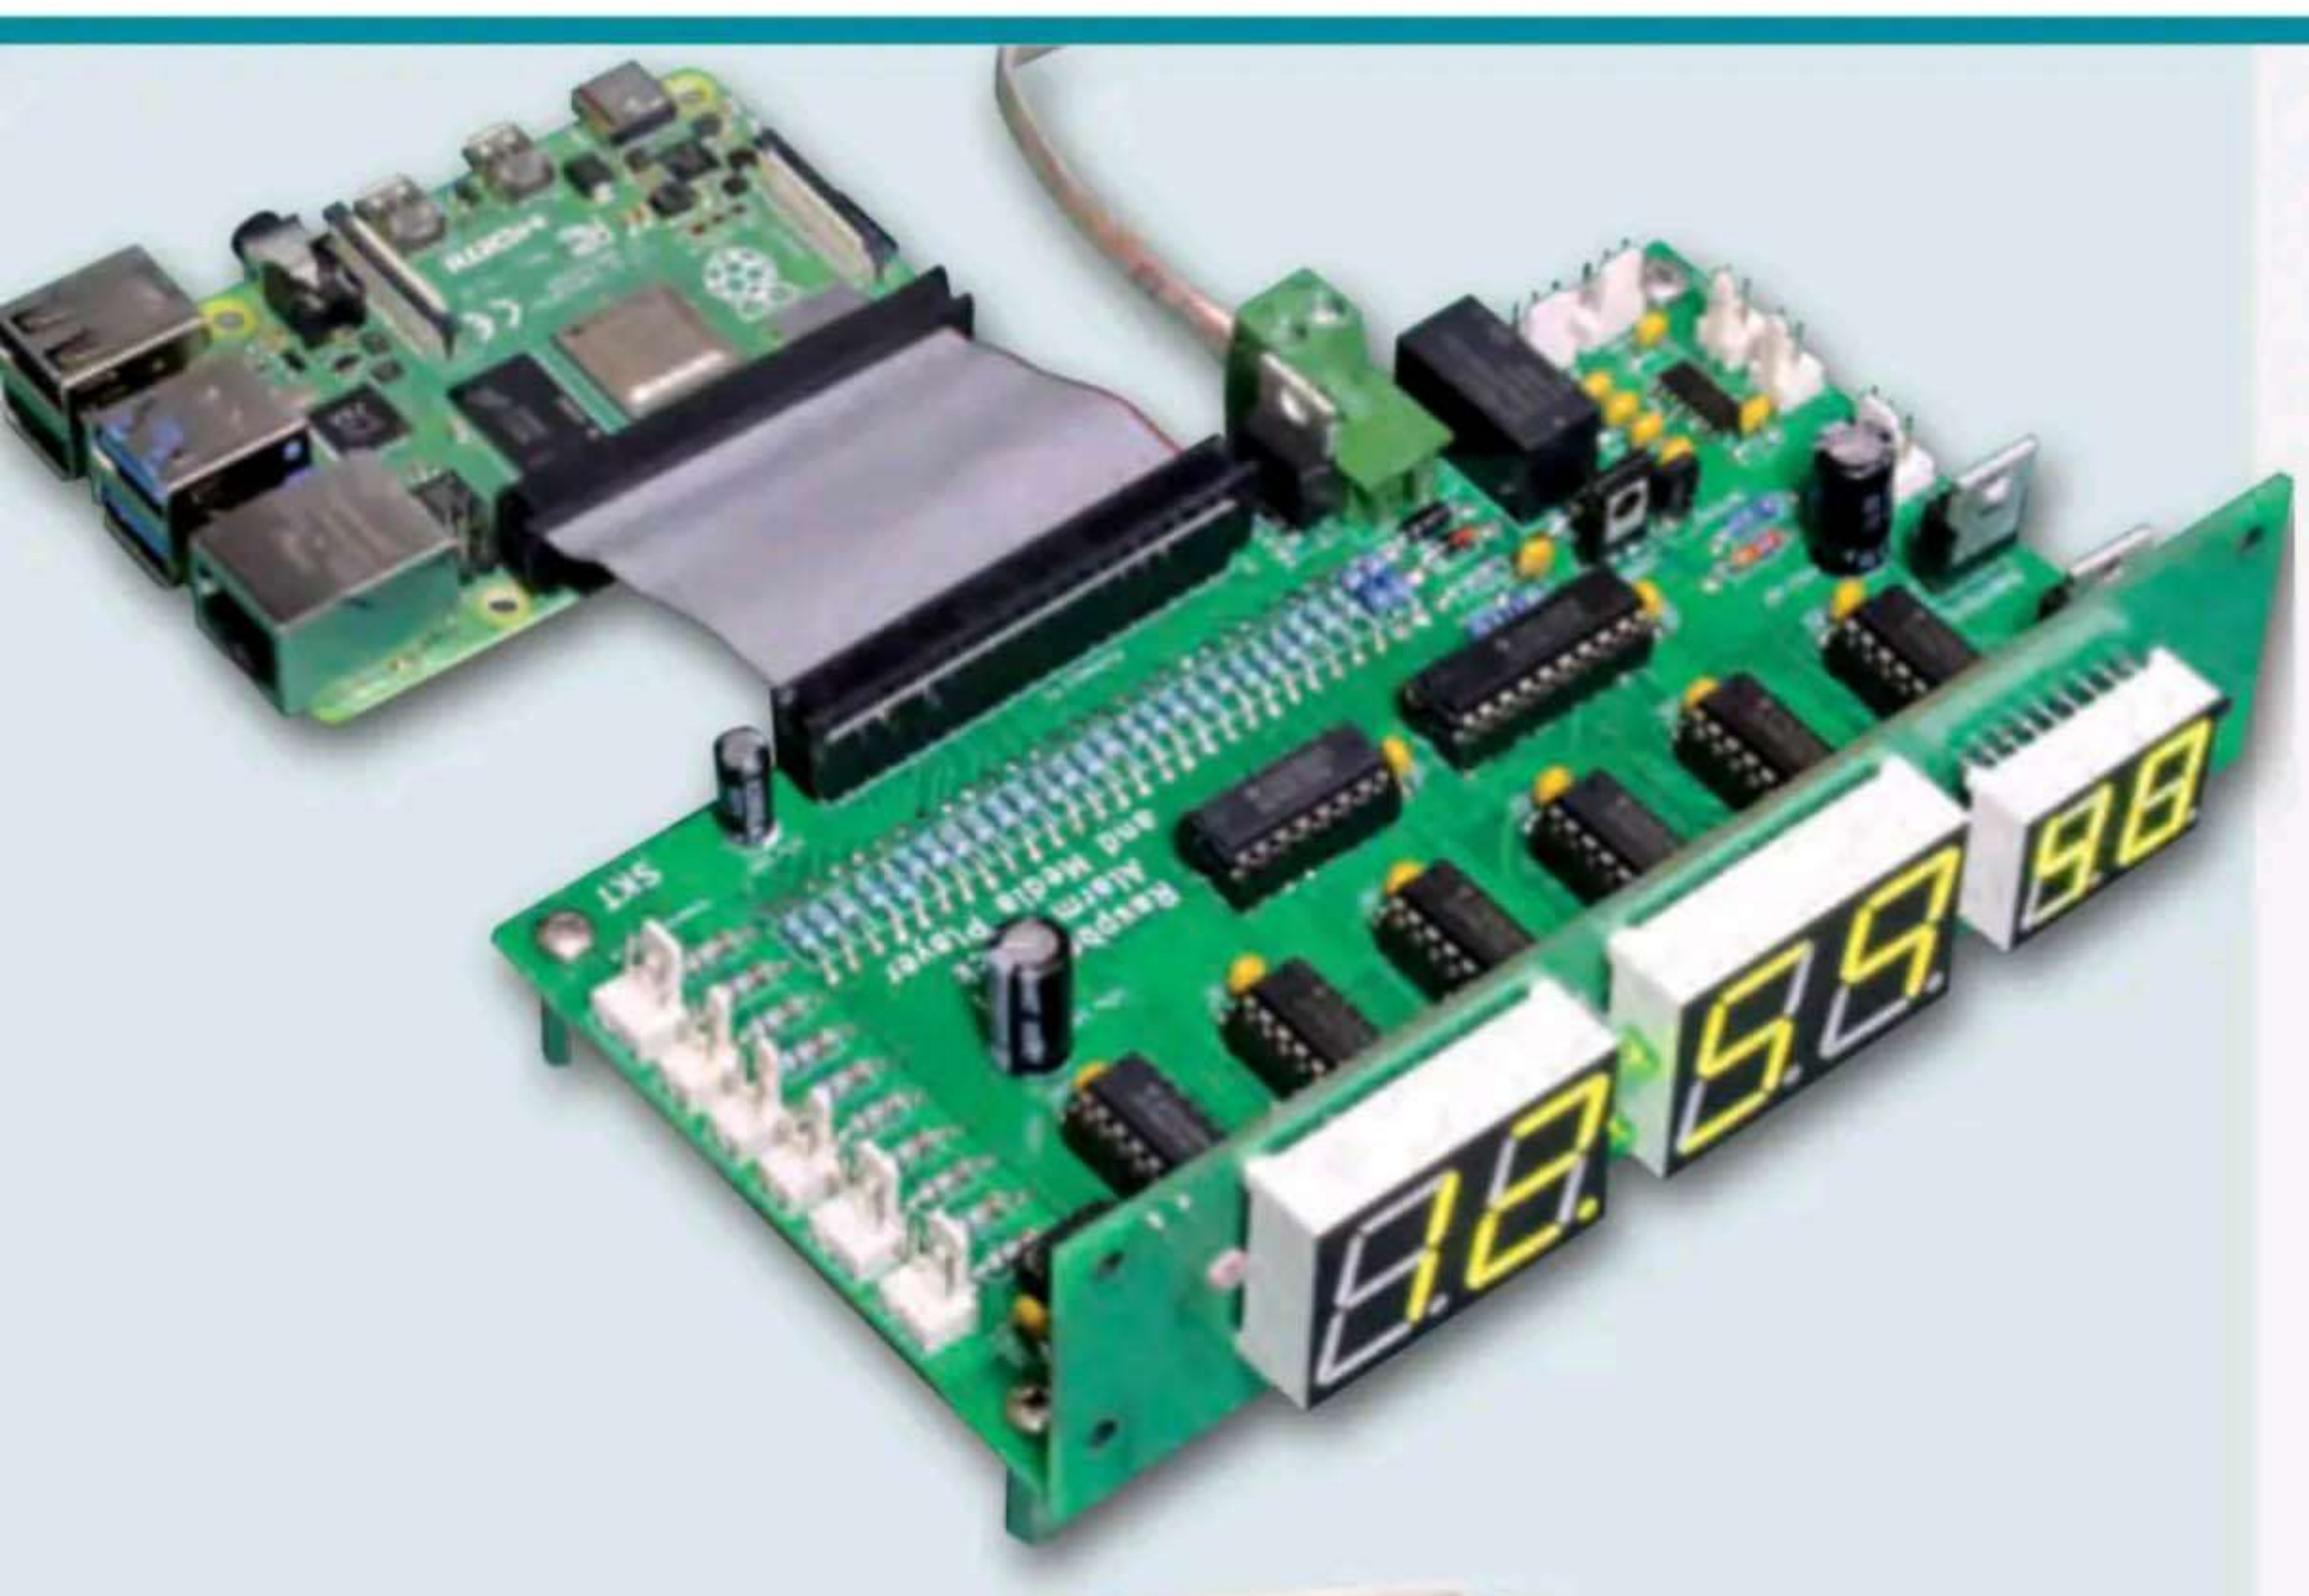

SECURE REMOTE SWITCH RECEIVER
Power
Relay
(Ack.)
Local

Projects and Circuits

Raspberry Pi Clock Radio by Stefan Keller-Tuberg 8
A modern alarm clock can take the trouble out of alarm setting and time keeping. This clock runs off most of the common models of Raspberry Pi (3, 4, Zero 2W etc) and can act as a Bluetooth media player.

The VGA PicoMite by Geoff Graham & Peter Mather 26
This easy-to-build, low-cost BASIC computer is based on a Raspberry Pi Pico. It has a colour VGA display output, PS/2 keyboard input and an onboard socket for SD or microSD cards. It's a great way to learn how to program or just have fun!

Secure Remote Switch, part one by John Clarke 54
Our Secure Remote Switch is designed to control low-voltage appliances, such as a garage door controller, fans, pumps or LED lighting. It's powered from 12V or 24V DC and can be controlled by up to 16 transmitters.

Multi-channel Volume Control, part two by Tim Blythman 66
Control the volume of up to 20 channels using this Volume Control system. This follow-up article covers the construction process and getting it running.

Series, Features and Columns

The Fox Report by Barry Fox 3
Overamplification ruins the sound of live performances

Net Work by Alan Winstanley 4
Continuing our 60th anniversary celebration with a look at the life of F.J. Camm, plus some of the pitfalls of home power backup for IT equipment.

Max's Cool Beans by Max The Magnificent 17
Arduino Bootcamp, Part 24 – it's time to get serial

6GHz Digital Attenuator review by Jim Rowe 33
It's inexpensive and adjustable from 0dB to 31.75dB attenuation in 0.25dB steps.

All About Capacitors by Nicholas Vinen 36
There are many different types of capacitor and each has its own set of pros and cons. This article has the information you need to choose the right type!

Techno Talk by Max The Magnificent 48
Tomorrow's World, today

Circuit Surgery by Ian Bell 49
An introduction to sample-based digital filtering

ZPB30A1 60W Programmable DC load review by Jim Rowe 62
This programmable constant-current DC load can be used for testing power supplies or checking the capacity of batteries. It is self-contained and good value for money.

Regulars and Services

Editorial 2
Coming up next year

Reader services – Editorial and Advertising Departments 2

Exclusive Microchip reader offer 25
Win a Microchip dsPIC33C Digital Power Starter Kit

Subscribe to Practical Electronics and save money 61

Practical Electronics back issue downloads 75

Practical Electronics PCB Service 76
PCBs for Practical Electronics projects

Classified ads and Advertiser index 79

Next month! – highlights of our next issue of Practical Electronics 80

Made in Australia & the UK.

Written in Britain, Australia, the USA and Ireland.

Read everywhere.

© Silicon Chip Publications Pty Ltd 2024

Copyright in all drawings, photographs, articles, technical designs, software and intellectual property published in *Practical Electronics* is fully protected, and reproduction or imitation in whole or in part are expressly forbidden.

The January 2025 issue of *Practical Electronics* will be published on Thursday, 5th of December 2024 (see p80).

Renewals paid by PayPal, Visa and Mastercard are now handled via our website: www.electronpublishing.com

Click the Subscriptions link at the top of the page. Select Publishing Services will continue to handle renewals via phone, cheque and mail.

Practical
Electronics

Volume 53. No. 12
December 2024
ISSN 2632 573X

Editorial

Editorial offices
Electron Publishing (Aust.) **Tel** +61 2 9939 3295
Unit 1, 234 Harbord Road **Email** pe@pemag.au
Brookvale NSW 2100 **Web** www.electronpublishing.com
Australia

Address mail to:
Electron Publishing (Australia)
PO Box 194, Matraville NSW 2036
Australia

Advertising enquiries
+61 2 9939 3295 pe@pemag.au

Editor	Nicholas Vinen
Publisher	Nicholas Vinen
Digital subscriptions	Stewart Kearn **Tel** 01202 880299
Online Editor	Alan Winstanley
Web Systems	Kris Thain
Production	Bao Smith
Technical staff	Tim Blythman, John Clarke

Print subscriptions
Practical Electronics Subscriptions
PO Box 6337
Bournemouth BH1 9EH **Tel** 01202 087631
United Kingdom **Email** pesubs@selectps.com

Technical enquiries
We regret that technical enquiries cannot be answered over the telephone. We are unable to offer any advice on the use, purchase, repair or modification of commercial equipment or the incorporation or modification of designs published in the magazine.

Questions about articles or projects should be sent to the editor by email: pe@pemag.au

Projects and circuits
All reasonable precautions are taken to ensure that the advice and data given to readers is reliable. We cannot, however, guarantee it and we cannot accept legal responsibility for it.

Some projects and circuits published in *Practical Electronics* employ voltages that can be lethal. Do not build, test, modify or fix any mains-powered equipment unless you fully understand the safety aspects involved and you use an RCD (GFCI) adaptor.

Component supplies
Silicon Chip Publications may offer kits or other parts for making our projects, but not in all cases. When kits are not available, readers will need to find and source parts themselves. We advise readers to check that all parts are still available before commencing any project in a back-dated issue.

Advertisements
Although the proprietors and staff of *Practical Electronics* take reasonable precautions to protect the interests of readers by ensuring as far as practicable that advertisements are bona fide, the magazine and its publishers cannot give any undertakings in respect of statements or claims made by advertisers, whether these advertisements are printed as part of the magazine, or in inserts. The Publishers regret that under no circumstances will the magazine accept liability for non-receipt of goods ordered, or for late delivery, or for faults in manufacture.

Transmitters/bugs/telephone equipment
We advise readers that certain items of radio transmitting and telephone equipment which may be advertised in our pages cannot be legally used in the UK. Readers should check the law before buying any transmitting or telephone equipment, as a fine, confiscation of equipment and/or imprisonment can result from illegal use or ownership. The laws vary from country to country; readers should check local laws.

Coming up next year

I mentioned in the November editorial that we had no Audio Out column in that issue because Jake Rothman was a little too busy. I had hoped that he would be finished moving house by now and would be able to contribute a column (even a short one) for this issue, but he unfortunately wasn't able to.

I expect that Audio Out will be back soon, but I guess we'll just have to wait until he is ready.

Teach-In 2024 has also run its course now. Mike Tooley is keen to start another series, but he needs a little while to prepare it. I expect we will publish the first part just after the middle of next year.

That leaves us with a bit of a vacuum in terms of articles from UK authors over the next couple of issues. We still have plenty of great content, it will just have to come mainly from our Silicon Chip writers for a few issues.

I have been looking at *PE* columns that were popular in the past and have contacted some of the authors of those columns. I am pleased to say that at least one has indicated he wants to return to writing for *PE*. However, once again, that won't happen straight away. Based on the author's statements, I'm hoping that column will also return by around the middle of next year.

This issue includes a fairly substantial article I wrote about capacitors. You might think that is a boring topic, but I feel that it turned out quite well, with a lot of useful information and helpful diagrams. We also have one more project article than usual this time.

In other magazine-related news, Clive Maxfield has decided that his *Arduino Bootcamp* series has run its course after two years. He will continue to write under the heading of *Cool Beans*, but it will have a slightly different focus starting next month. I'm anticipating that his new series will retain the 'hands on' aspects of *Cool Beans* but with renewed vigour and focus.

Nicholas Vinen, Electron Publishing (Australia)*
Publisher & Editor,
Practical Electronics Magazine

* a division of Silicon Chip Publications Pty Ltd.

The Fox Report

Barry Fox's technology column

Overamplification ruins the sound

For many years now, I have been sounding off about electronic amplification for live music; specifically, the mis-use of electronics.

The generally accepted best definition of 'hi-fi' is the 'closest approach to the original sound'. But how is it possible to judge electronic reproduction of musical instruments without knowing what live music sounds like?

So I go to a lot of live music events. Classical music – solo piano, small group and full orchestra – remains mainly unamplified. Overloud amplification, and deliberate 'fuzz' distortion, has become an integral part of rock and pop. Jazz (including folk and blues) is where amplification is a choice. And all too often, live jazz is badly amplified. So the live reference point is distortion, mostly caused by system overload.

The root of much amplification evil is the overlooked fact that most musical instruments, professionally played, will generate a higher sound level (100dB or more at a few metres) than a reasonably affordable stereo amplifier, unless it is turned up to a volume setting which distorts and clips the output signal.

Affordable speakers will add their own distortion. Add to that the ever-present risk of feedback, with sound from a speaker getting picked up by a microphone to create howlaround. So the net result of amplification is all too often to make live music sound worse than a home hi-fi.

Remember that performing musicians have no idea how the audience is hearing them; they rely on on-stage monitors or sometimes earpieces. If the speakers that fire into the audience are distorting, no-one on stage will know.

Hall acoustics also play a big role in what the the audiences hears. For many years, London's Albert Hall was infamous for the long echo caused by the high dome ceiling. Composers joked about an Albert premiere being the sure way to get a new work heard twice – once live and again a few seconds later!

The Albert Hall echo was finally tamed by suspending 'flying saucer' sound absorbers high above the audience. Churches with high ceilings may be good for choirs and rolling thunder organs but they can be a bad place to hear most other music. I once heard the fine singer Willard White destroyed by the Alpine echo of a North London church.

The world-famous Studio One at the Abbey Road Studios was built with seaweed sound-damping cladding on the walls, which could be removed to make the sound more live. It has long since been replaced by fire-safe foamy material, of course.

The Royal Festival Hall on London's South Bank was for years the opposite – 'dry' with inadequate echo – and relied on low level all-round electronic amplification.

In the early days of jazz, in New Orleans and Chicago, there was no electronic amplification. Singers used mechanical megaphones and the bass came from wind instruments like tubas. When electronic microphones and valve amplifiers first arrived in the 1930s, they were used mainly by vocalists and for announcements.

Listen to live recordings from that era and you will hear very little piano and next to no bass or rhythm guitar. One of the best examples is the seminal recording of the concert in Carnegie Hall by the Benny Goodman big band in January 1938.

Although it became one of the best-selling recordings ever, and one of the first vinyl LPs, its commercial release was never planned. Goodman's friend Albert Marx set up three microphones and cut lacquer discs that Goodman then forgot about for 12 years, before finally listening to them and authorising their sale in 1950.

Benny Goodman never did learn how to handle amplification. Loren Schoenberg, a respected saxophone player, music historian and the driving force behind the National Jazz Museum in Harlem, organised and played in Benny Goodman's last band from 1980 to 1985 (while archiving his memorabilia and music before they were lodged with Yale University). When I interviewed Schoenberg, he recalled:

Pete Long (clarinet), Anthony Kerr (vibraphone), Colin Goode (piano), Jo Pettitt (bass) and Bobby Worth (drums) recreating no-amp 1930s/1940s music at a venue in Sussex.

Benny would arrive for an open air concert where the sound engineers had carefully set up a dozen microphones. He'd wave his hand saying "I don't want any of that" and then play without any amplification so that only a few people could hear him.

I can personally vouch for that. When the "King of Swing" performed at the 1982 Knebworth open air festival, he ostentatiously ignored all the microphones on stage. The result was that only the first few rows of the audience heard anything of his playing.

The 1938 Carnegie concert has recently been recreated live, with almost no amplification, by British bandleader/clarinettist Pete Long and drummer Richard Pite, for instance in London's Cadogan Hall, which has excellent live accoustics. Long is now involved in similar no-amp re-creations of Benny Goodman's smaller groups. Bass players in particular are now so used to relying on electronic assistance that only a few acoustic string bassists can handle the recreation job.

If you like that kind of music and care about hi-fi, watch out for this type of recreation event. Who knows, the idea may catch on wider. The pop and rock world is already drowning in tribute bands recreating famous names.

A final thought: there is a new and rather annoying industry based on sending out questionnaires to customers immediately after they have bought anything or been to any event. Avoid the temptation to simply bin or delete the question form; instead, grab the opportunity to tell event organisers when and why the sound you paid to hear was bad.

Net Work

Alan Winstanley

This month's feature continues our 60th anniversary celebration with a historian's look at the life of F.J. Camm. Also, one year after upgrading his home network power supply protection, the author brings disappointing news of failure!

I hope that readers enjoyed last month's retrospective feature celebrating 60 years of *Practical Electronics* magazine. I described how a range of popular 'Practical' magazines sprang up to meet the needs of post-war householders, mechanics and hobbyists who eagerly tackled hands-on projects or home repairs with great enthusiasm.

The magazines owed their existence to the unswerving dedication and the devotion of one particular Editor, the legendary Frederick James Camm, who worked for George Newnes.

I'm delighted to say that I have managed to obtain a copy of a rare book entitled "F.J. Camm, The Practical Man 1895–1959" written in 1996 by the late Gordon G. Cullingham. My second edition reprint is autographed and numbered 38 of a 100-copy print run, and it will undoubtedly have pride of place in my collection.

Mr. Cullingham was a respected local historian and highly-regarded president of the Windsor Local History Publications Group, Windsor being the English city where F.J. Camm was raised.

Mr. Cullingham researched the life and works of 'F.J.' in great depth, and his appraisal offers a detailed insight into what made F.J. tick, and how his passion for practical pastimes evolved, starting from his boyhood.

Readers are regaled with many fascinating facts about F.J. that have probably passed unnoticed in our own hobby electronics arena. We are told that his grandparents were actually named "Cam", but misspellings on some Royal Association certificates saw the surname "Camm" printed instead. The family liked it so much that the name stuck.

After reading Mr. Cullingham's book, I suspect that there is probably a little bit of "Camm" in all of us, for F.J. exemplified those traditional British qualities of ingenuity, industriousness and inventiveness.

We learn how, as a young boy, he became interested in aircraft and the prospect of manned flight in machines that were heavier than air. A local talk brought news of a historic first 'full circle' powered flight in 1904 by the Wright brothers in the USA.

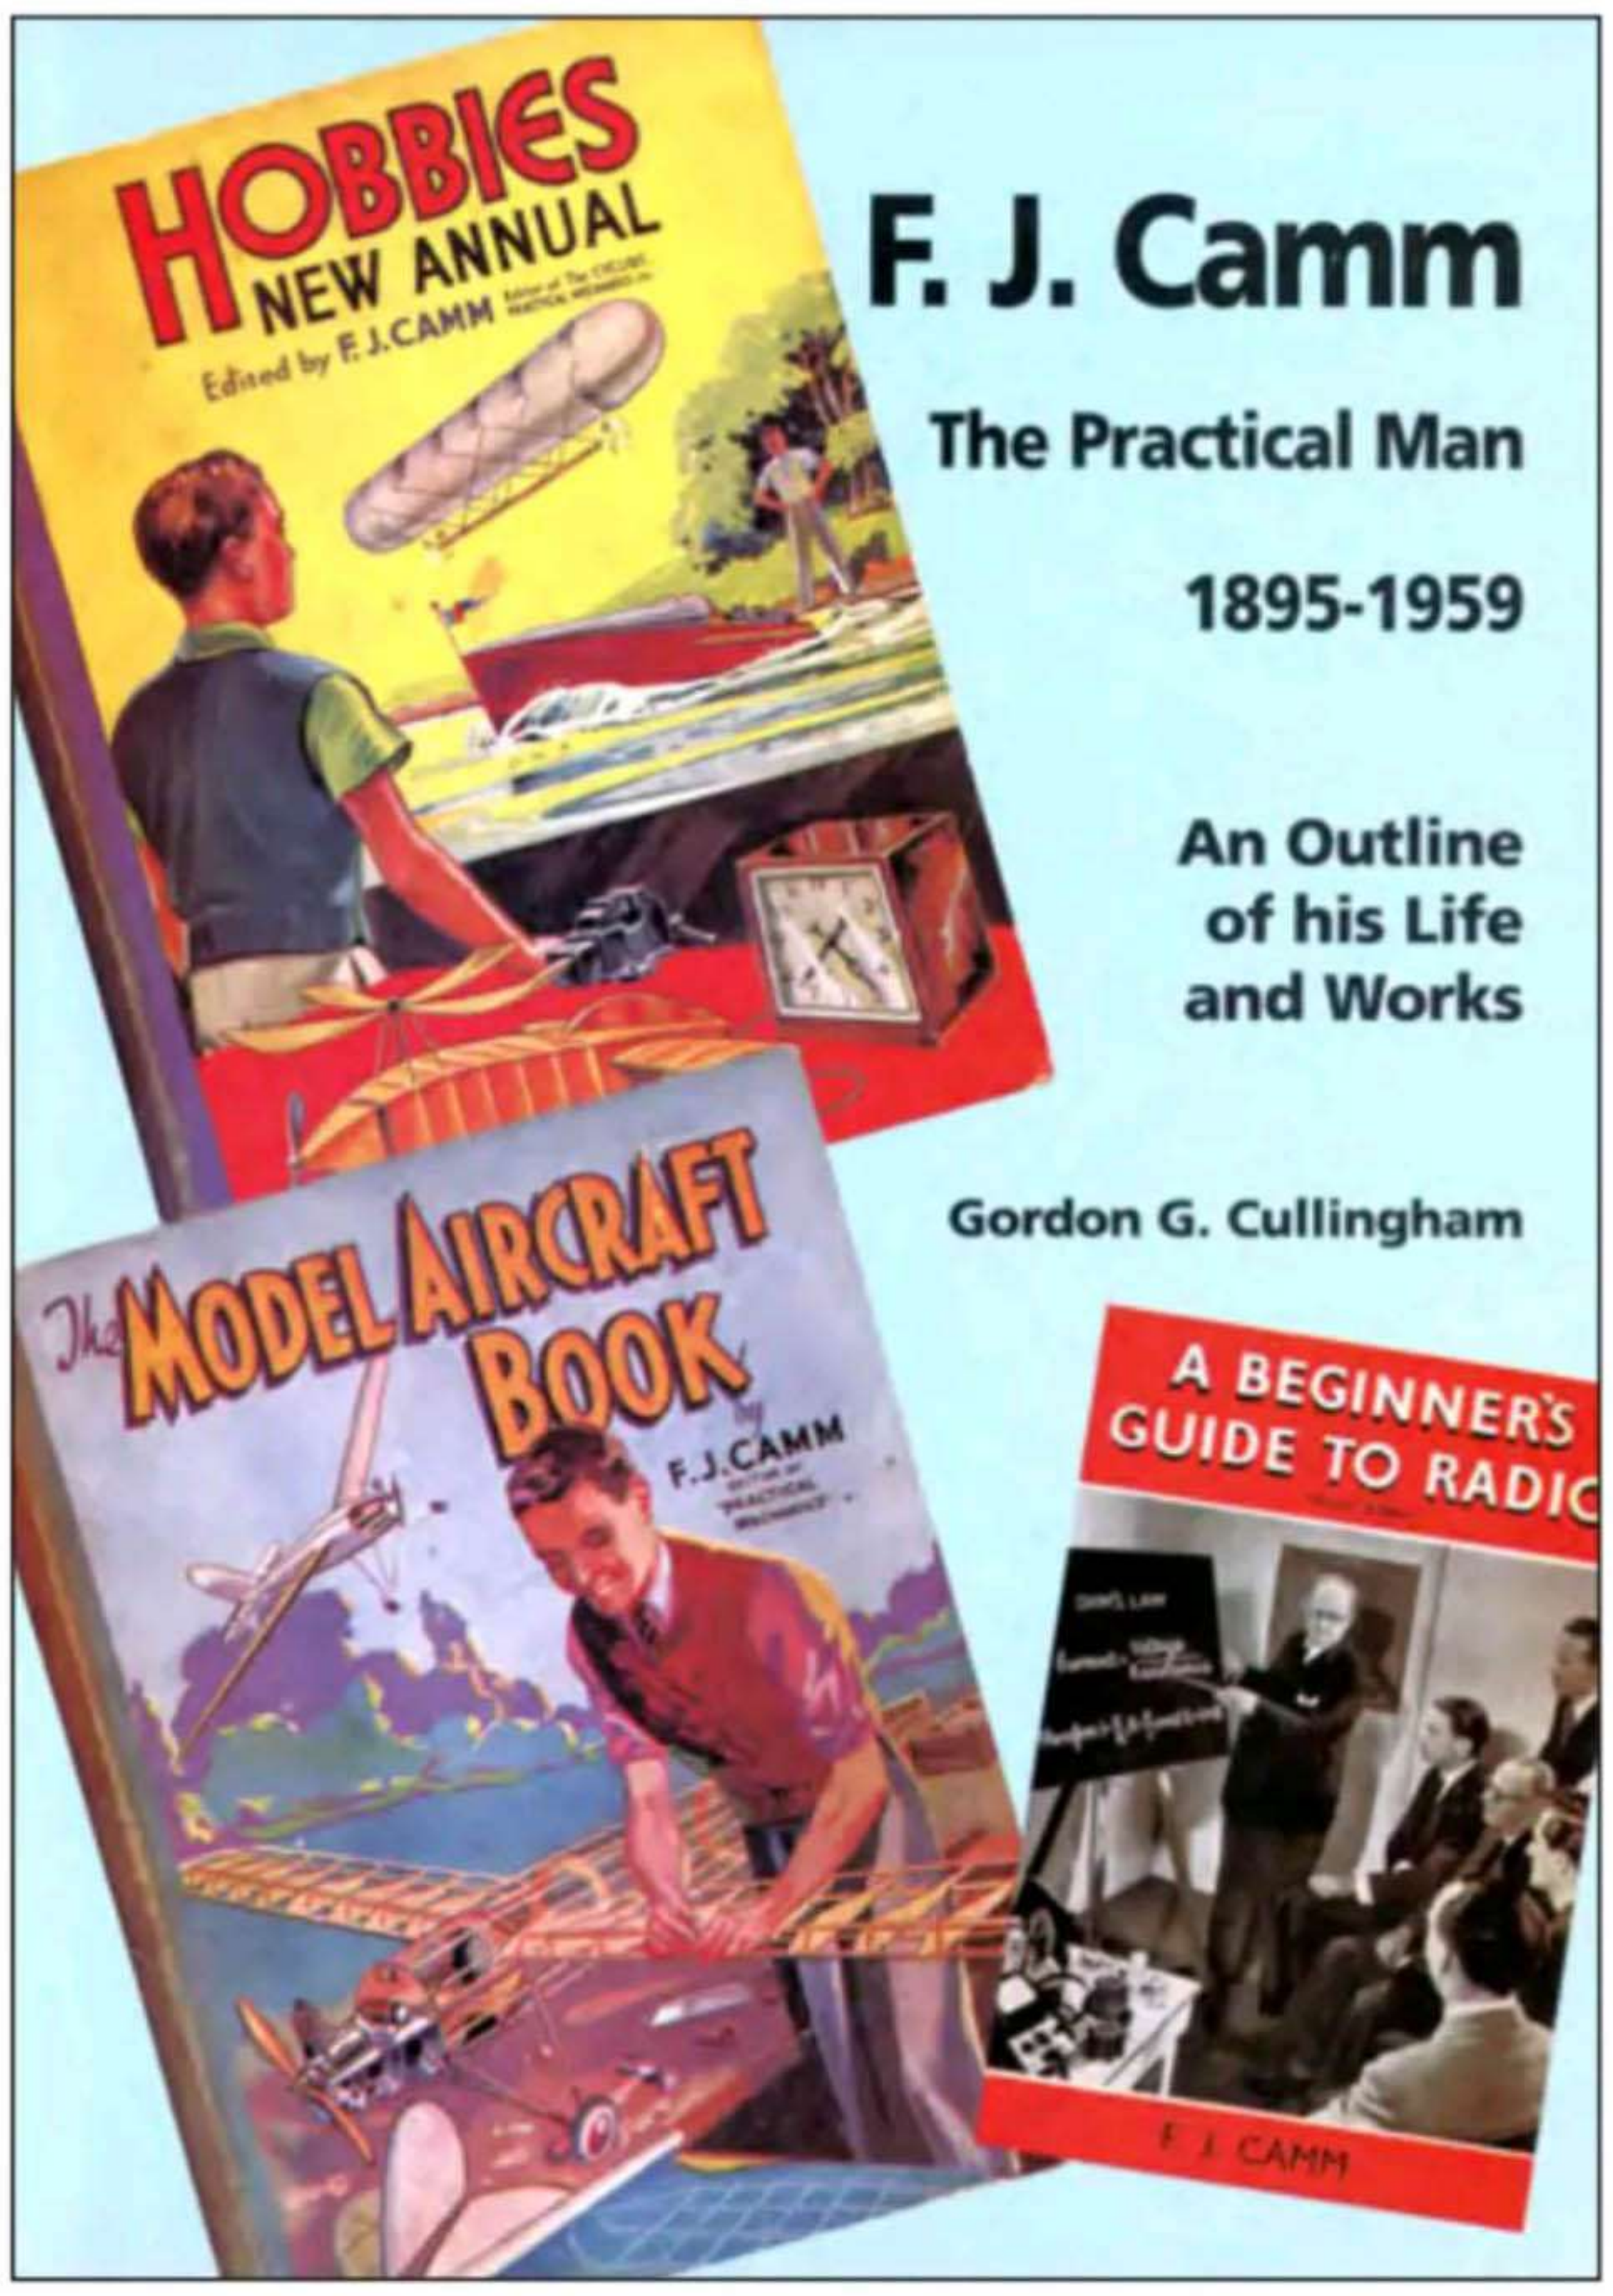

An appraisal of F.J. Camm's life, privately published by the late Gordon G. Cullingham of Windsor.

According to the book, the talk held in Windsor was presented by Patrick Alexander, a distinguished English aviation enthusiast and balloonist who knew the Wright brothers personally. His own fascinating life is portrayed at **https://en.wikipedia.org/wiki/Patrick_Young_Alexander**

As the Camm family lived in Windsor, on the same road as the talk's venue, the historian speculated that some of the exciting news of the Wright brothers' flight probably reached the ears of a 10-year-old F.J. Camm.

Hurricane force

It also seems likely that F.J. Camm's elder brother, Sydney, was captivated by early drawings of the Wright brother's biplane. Sydney Camm would eventually go on to design the Hawker Hurricane, the formidable fighter aircraft that really won the Battle of Britain, although the more glamorous Spitfire received most of the credit.

Wikipedia has more on Sir Sydney Camm at **https://en.wikipedia.org/wiki/Sydney_Camm** and a blue commemorative plaque appears on the Camm family home, a non-descript row of Victorian terraces in Windsor that can be seen today at **https://maps.app.goo.gl/sFChp6EA8qaWgmRr6**

At the start of the 20th century, not only were early radio and electromechanical devices capturing the hobbyist's imagination, but airships and aeronautics were as well. Plans and kits for model airplanes were being snapped up by those keen to harness the principles of flight, albeit on a modest scale.

Mr. Cullingham's book explains in great detail Camm's fascination with model gliders, airships and aircraft of the era. He would become a gifted designer and illustrator; it is thought that he introduced expensive balsa wood (actually classed as a hardwood) from the USA for the first time, which greatly facilitated the construction of lightweight model gliders.

F.J. also wrote numerous handbooks and columns in periodicals in the 1920s for publishers such as Cassells, Pitmans and George Newnes, on topics like model aircraft, propeller design, mechanical projects and much more besides.

The advent of Morse code wireless transmission doubtless piqued his interest too, notably when the Titanic was lost in 1912.

A Marconi telegraph boy named Owen Allum had perished on the Titanic. He was originally from Windsor and, highly unusually after the catastrophe, his remains were duly repatriated to England, so "all Windsor knew of the tragedy", as George Cullingham wrote.

There is more background on this at **https://windsorlocalhistorygroup.org/drowned-at-sea-owen-george-allum/**

F.J. Camm would eventually have to learn Morse code when applying for a radio license anyway, which probably

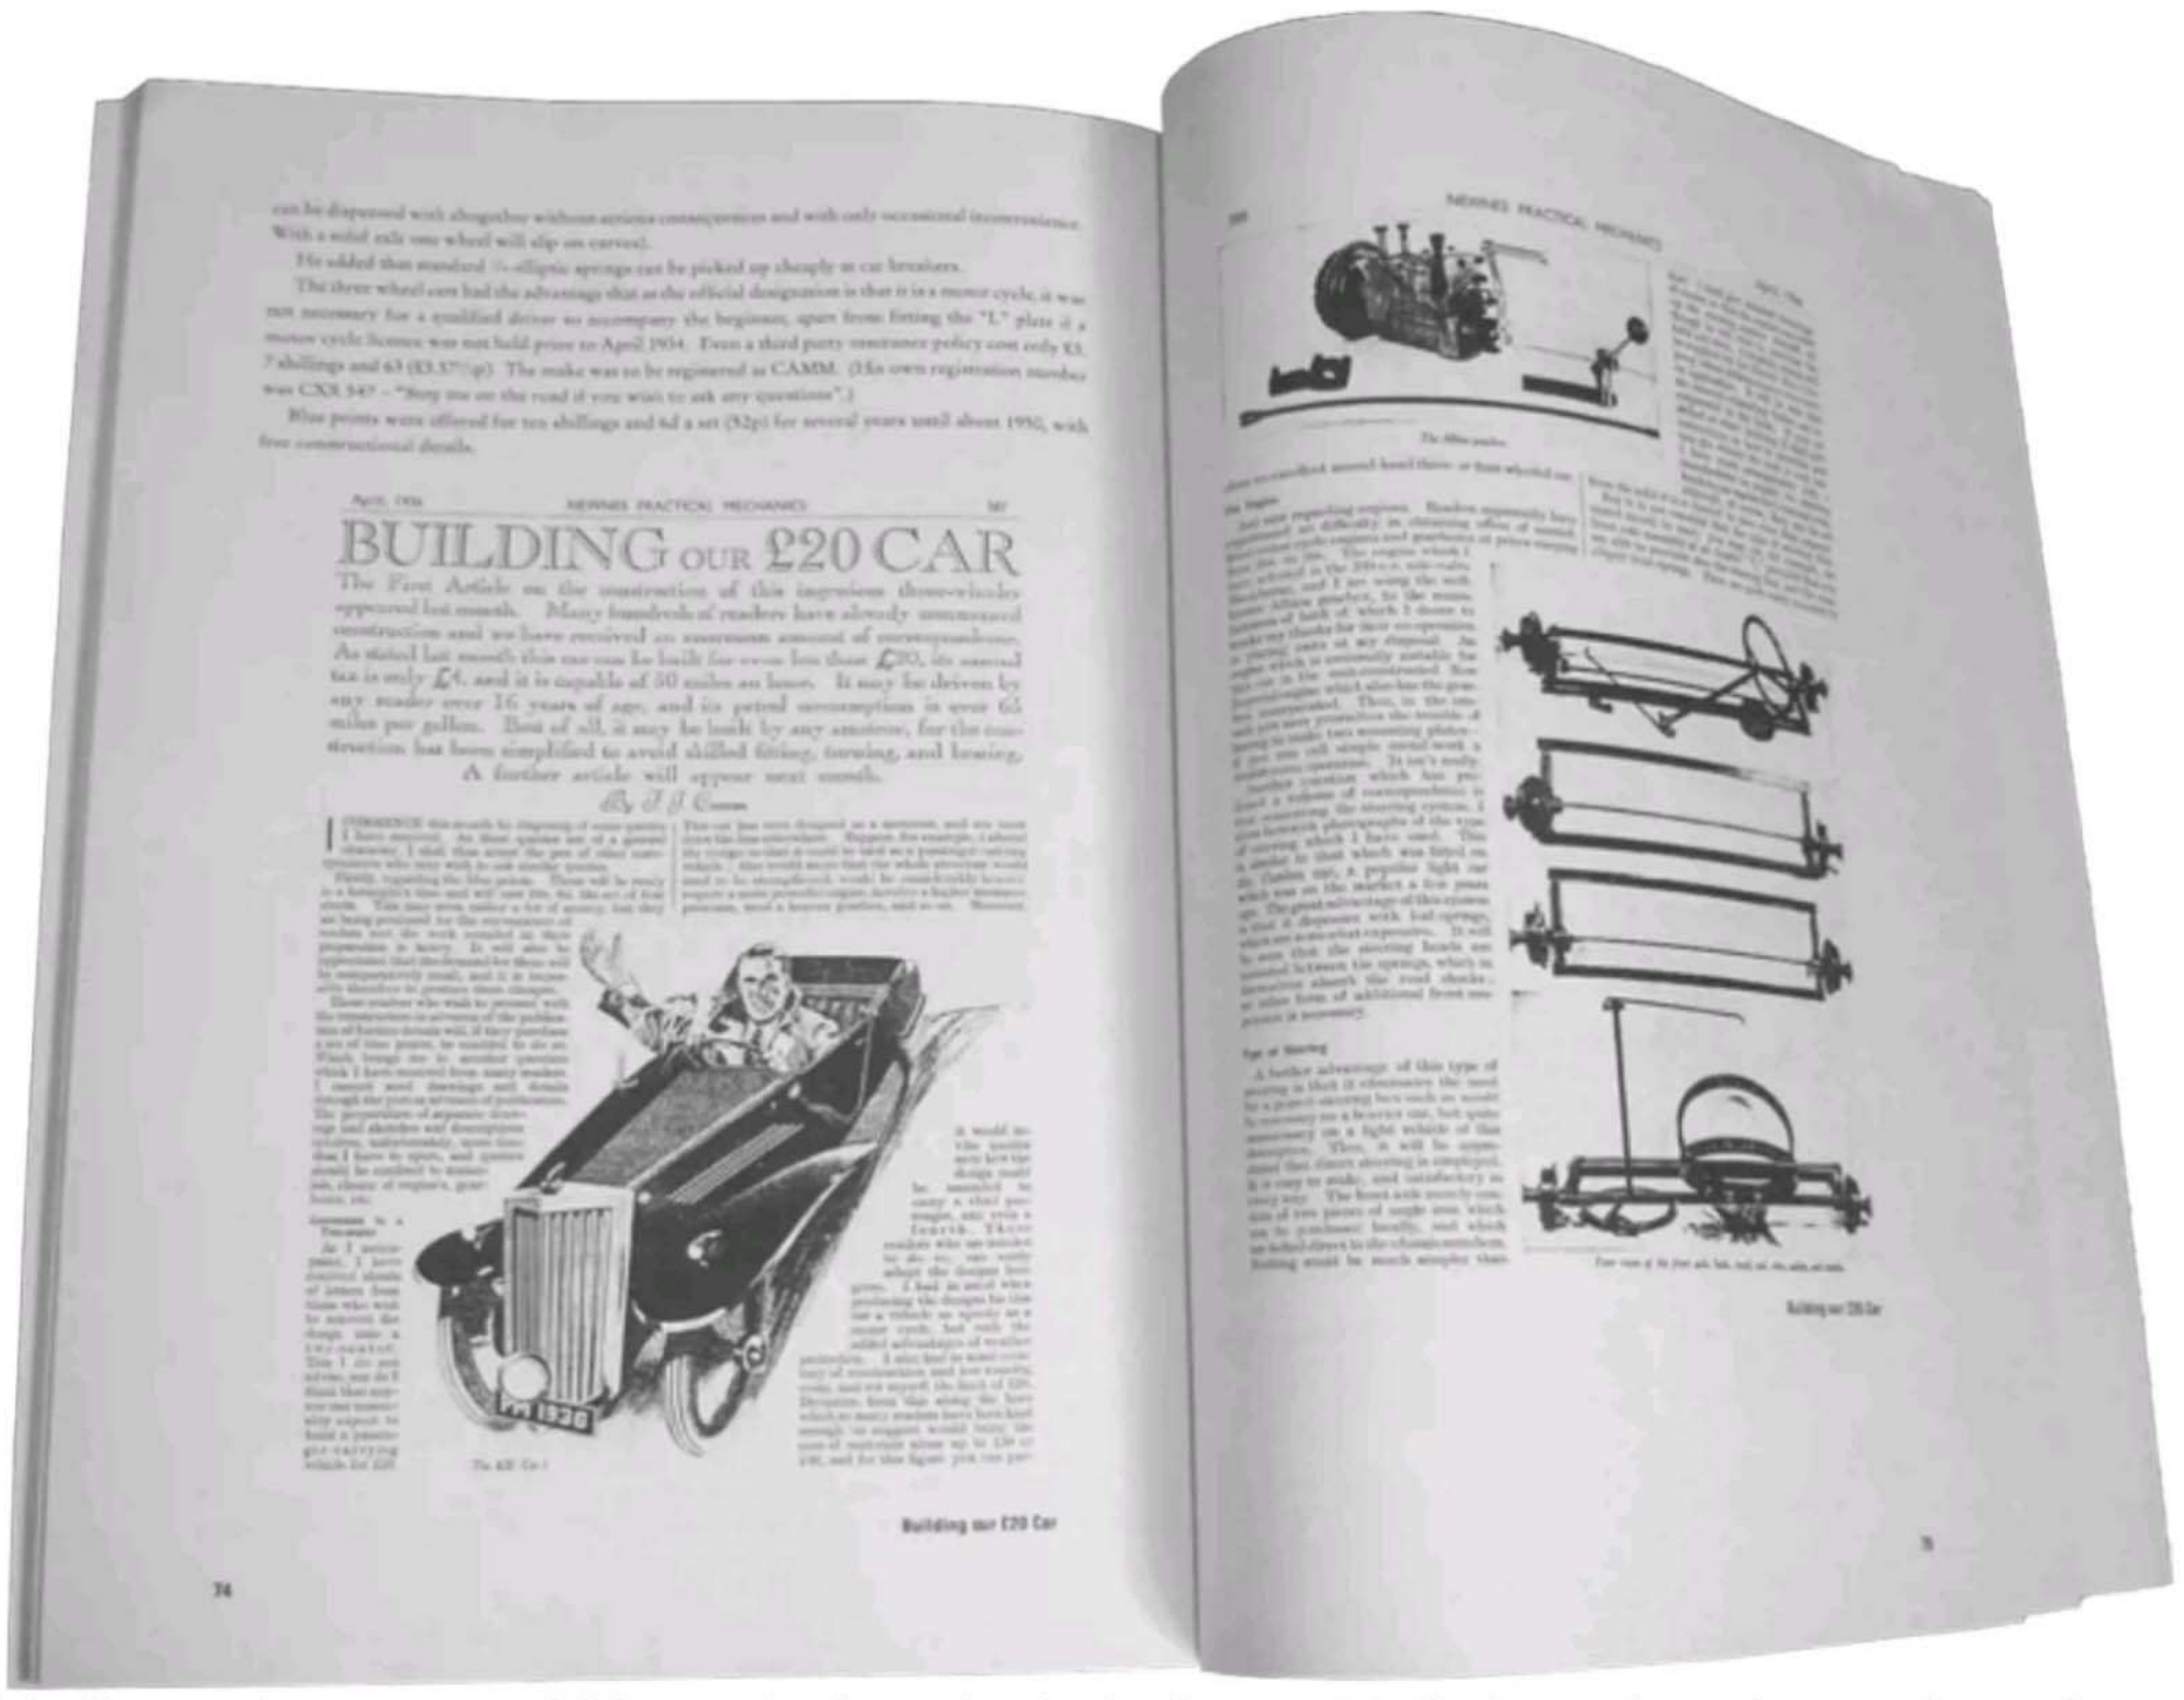
BUILDING OUR £20 CAR

F.J. Camm's more ambitious designs included a 1936 timber-chassis car that almost anyone could build for under £20, or so he claimed.

coincided with his interest in early radio design.

In Newnes' Practical Mechanics of April 1936, he offered a design for a "£20 car", a single-seat three-wheeler contraption with plywood chassis, powered by a 500cc motorcycle engine that "can be built by any amateur with limited tools". A four-wheel, two-seater 'Midget' car followed in 1937, and his practically-minded followers had no problem puzzling out adaptations for brakes, suspension, axles and the like.

Gradually, all these elements of hands-on engineering, theory, inventiveness, design, drawing and publishing skills gained in aeronautics, mechanics and radio would form the core of F.J. Camm's unrivalled portfolio of practical expertise and know-how.

Camm was also highly defensive of his work, and he did not shy away from using brutal litigation to uphold his reputation. His work was sometimes pirated and the outcome of several legal cases are described in the book.

A heavy smoker, he died from bronchial pneumonia in 1959, aged 63. Newnes itself eventually became an imprint of Butterworth-Heinemann, publishers whom I met in Oxford in the 1990s, and Newnes 'Pocket' books became a standard fixture on many engineers' and hobbyists' bookshelves and work benches at the time.

Punctilious punctuation

I described last month how, in 1964, *Practical Electronics* was spun out of the long-established title *Practical Wireless*. In December 1977's *PE*, the magazine's founding Editor Fred Bennett bade farewell to *PE* as the magazine was relocating from London to Poole in Dorset. Both he and Dave Barrington, the production editor, would stay behind in London to work on sister title *Everyday Electronics* instead.

Dave, who started his career working for Newnes as a 15-year old 'copy boy' (an inter-office errand boy), was especially saddened by the move, and he wrote in his final 'Market Place' column that he was the sole remaining editorial staff member of the Practical group of magazines who had worked with, and was trained by, F.J. Camm himself.

In the years that followed, I would often have a discussion with Dave over an errant full-stop or apostrophe that his eagle-eye had spotted in my own material, and so I gradually refined my manuscripts mindful of such minutiae. I am sure that Dave's quest for correctness and accuracy was a direct result of F.J. Camm's influence.

This meticulously-researched history book includes reprints and snapshots of magazine milestones too numerous to cite here, but F.J. was undoubtedly the father of technical book and 'practical' magazine publishing in Britain, and his influence was profound and enduring.

If you would like to learn more about his life's work, F.J. Camm, *The Practical Man* 1895-1959 (A4 format, 100pp, ISBN 0 952 8448 0 X) is a very thorough and unmatched resource.

It was privately published, but copies of the limited 100-run reprint occasionally appear online at AbeBooks, eBay and Amazon. It's worth noting that prices vary considerably, and mint condition copies sell at a premium.

The Windsor Local History Group provides considerably more information about the late George G. Cullingham on their website at the URL: **https://pemag.au/link/ac1s**

Feeling powerless

This morning's email started with one generated by my Synology NAS, complaining once again about an "improper shutdown".

A year ago, in the December 2023 issue, I described some options for small uninterruptible power supplies (UPSs) that could, in the event of a power glitch, safeguard typical home IT equipment such as routers or modems.

My main PC and dual monitors already utilise a bulky Salicru-brand fan-cooled 1.2kW UPS containing 24V sealed-lead acid batteries. It allows plenty of time for files or programs to be shut down gracefully, hopefully averting the risk of file or disk corruption if the power fails unexpectedly.

A typical UPS has multiple IEC-style mains outlets to power a few small peripherals, but I'm not keen on the idea of leaving a hefty UPS running unattended continuously in a semi-domestic environment. Some compact DC supplies are sold that contain a Li-ion battery instead, which I felt would be more appropriate for protecting small network peripherals like my Synology NAS, modem, router and five-port switch.

The iPower-H by Power Inspired (**https://www.powerinspired.com**) is a device I described last December. It is

NEWNES
Electronics Toolkit
Farnell
GEOFF
NEWNES
MS-DOS
POCKET
NEWNES
ELECTRONICS
ENGINEER'S
POCKET BOOK
Keith Brindley
A Butterworth-Heinemann Book

Newnes Pocket Books were a common sight on hobbyists' and engineers' bookshelves to the 1990s and beyond.

The author's two remaining compact UPSs are used to keep the network running during a power cut. They are compatible with equipment powered by 9V, 12V or 24V DC.

a mains adaptor style UPS containing a 3.6V 18650 battery that supplies 12V in the event of a mains outage (see *Net Work*, December 2023).

Alternatively, the iPower-DC2 is a paperback book-sized UPS with a Li-ion battery that fits in between the equipment and its mains adaptor. It has five DC outlets that can be set for 9V, 12V or 24V (to match the voltage of the ITE's mains adaptor). You can expect up to 20-30 minutes or more of protection, depending on the load.

I also wrote how I'd grappled with different DC plug sizes, adaptor leads and Y-splitters to accommodate a motley range of connectors. I found that most (but not all) gear uses 5.5mm outer diameter, 2.1mm inner diameter barrel plugs for low-voltage DC power these days. However, some barrel connections and adaptors fit together very poorly and can fall apart at a whim, causing more problems.

Having suffered the cost last year of restoring my Synology NAS following a string of power cuts, I configured it to remain powered down after a power failure. My cunning plan would safeguard my home network against any more outages, or so I thought...

I complained to Northern Powergrid, my electricity supplier, about the damage wreaked by power cuts and received, as one would expect, their boiler-plate response disclaiming all liability.

To add insult to injury, they added that "manufacturers should build their equipment to withstand power cuts", so it's all our fault anyway.

Several times during this year, I then received email alerts from my Synology noting that "Diskstation was shut down improperly" and when I checked, the disk drives had indeed powered off completely.

Oddly, though, the DC2 UPS that powered the NAS had switched itself off as well. I knew its battery was OK because I could switch it on again manually.

I took this up with the manufacturer, who agreed that the UPS should comfortably supply the 25W @ 12V needed by my Synology. My NAS only has two disks, whereas they used four drives needing twice the power in their testing.

A NASty problem

Eventually, the problem was blamed on the NAS itself. The theory was that an inrush surge possibly overloaded the UPS when the current was switched from 'straight through' mains power to battery power instead. The makers state that the UPS would shut down automatically if it sensed an overload.

Yet I couldn't replicate my problem by plugging or unplugging the NAS supply to feign a power cut; the Synology carried on running as it was supposed to.

With no solution or workaround available, I've written off that idea as something that should have worked 'on paper' but didn't in practice. The NAS is now mains-powered again.

My mood didn't improve when the 12V adaptor-style iPower-H used on my router started flashing angrily when I simulated another power outage. The router powered off and the UPS stopped working, too.

I discovered that the removable (once you know how) 18650 (18mm diameter, 65mm long) 3350mAh cell had failed after less than a year. My confidence dented, I abandoned that idea as well.

The two iPower-DC2 supplies that remain now power the 12V DC router and modem, with 9V DC powering a TP-Link mesh unit and a five-port switch.

I sourced another Zolt DC 34W mains adaptor, set for 12V, from Amazon as before (**https://pemag.au/link/ac1r**). This set-up seems to run happily and the network has kept running through several power cuts.

One problem I can't overcome is that my TP-Link Tapo smart bulbs light up once electricity is restored. Several times, they were found to have been blazing all night after a power cut!

Such uncertainty is another reason I won't ever trust any brand of smart socket to operate power-hungry appliances or whitegoods.

Next month, I'll take a nostalgic look at a British product that changed the hobby electronics scene for ever. It's the 'balsa wood' for electronics constructors, but what was it really?

As always, you can reach the author at **alan@epemag.net** PE

HAMMOND®

1550ZF IP68 flanged die-cast aluminium

Learn more: hammondmfg.com/1550zf

uksales@hammondmfg.com • + 44 1256 812812

PCBWay

TURNING IDEAS INTO REALITY

PCBWAY EMPOWERS INNOVATORS WORLDWIDE

Join the **PCBWay revolution** today and turn your **imagination** into **reality!**

WWW.PCBWAY.COM

SERVICE@PCBWAY.COM

ENGINEERED TO PERFECTION

Our state-of-the-art facilities and expert team of engineers ensure that your designs are brought to life with the utmost precision and efficiency, from a small-scale prototype to a large-scale production run.

QUALITY ASSURANCE

Illustrate PCBWay's rigorous quality control processes, emphasizing that every product undergoes meticulous testing to ensure it meets the highest industry standards.

GLOBAL REACH

Highlight PCBWay's extensive network, reaching innovators across the globe. Mention our fast and reliable international shipping options.

CUSTOMER SUPPORT

Emphasize PCBWay's commitment to outstanding customer service, featuring a dedicated team of experts available 24/7 to assist with any inquiries or support needs.

Raspberry Pi Clock Radio

Every day for over five decades, I awoke to the mono lo-fi sound of a clock radio. On the occasions I needed to be woken for an early flight, I worried about accidentally mixing up AM and PM when resetting the alarm, and sometimes that actually happened. Twice a year, my old clock needed daylight savings adjustments, and when the power went out, it flashed 12:00. I decided to fix all that!

Part 1 by Stefan Keller-Tuberg

A modern alarm clock can sound great, keep precise time and support multiple alarms. Alarm settings should include the day(s) of the week as well as the time, and you should be able to decide what sound each alarm plays, for how long, at what volume, and whether it's a one-off or will repeat indefinitely.

There's no longer any reason you should need to set the clock's time. It can be accurately fetched over the internet, with daylight savings and leap seconds adjustments applied automatically. Also, if the clock has speakers and a wireless connection, why not support streaming audio from a LAN, the internet or a smartphone?

This project is based around a Raspberry Pi and supports all of these 'dream' features and many more; it can even snooze or cancel the alarm on your partner's clock!

The Raspberry Pi is a great platform because many of the required capabilities are already built in. Also, many of us already have a Raspberry Pi or two gathering dust and waiting to be put to good use!

The custom hardware can interface with any Pi that runs Linux with a network connection and a 40-pin expansion connector. It has been tested to work with the Pi 3, Pi 4 and Pi Zero 2W models. It should also work with a Pi 2 or Pi Zero W, but they haven't been tested yet. The newly released Raspberry Pi 5 is not suitable as it lacks an analog audio output.

The Pi 3 and Pi 4 have the most capable processors and are therefore the best options, especially for media streaming. They also have Bluetooth, so you can stream audio to the Clock from a smartphone or similar. The slower Pi variants may be suitable if you plan to integrate a traditional over-the-airwaves radio as the audio source. In general, we recommend using a Pi 3 at the minimum.

Hardware features

The design uses two PCBs: a display board and a main board. The clock hardware and the Pi are powered from the same 5V DC source. The Pi connects to the main board and receives power via a short ribbon cable. I have used plugpacks for the four clocks I built for myself, but if you have the space, you could integrate a power supply inside the enclosure.

The display comprises large 20mm-tall hours and minutes digits with 15mm-tall seconds digits that will be prominent anywhere in the room. The minimum and maximum LED brightness range is configured via the clock's web interface, with the brightness automatically adjusted within the set range in response to changes in ambient light conditions.

At its brightest, the display can be read in a sunlit room; at its dimmest, it is unobtrusive at night.

The physical user interface comprises six switches: three centre-off toggle switches and three momentary contact push buttons. The switches replicate the features commonly found on clock radios, including snooze and media player controls, but there are too many features to control with six switches alone.

All features are accessible using a web browser, so you can control the clock from a computer, phone or tablet connected to WiFi.

A built in stereo amplifier with digital volume control can drive internal or external speakers according to your construction preferences. The Pi is the primary audio source, but you can integrate an alternative source, such as a traditional radio, if you want to.

While the Raspberry Pi analog audio is not quite hifi quality, it is not to be scoffed at. I used a pair of bookshelf speakers with one of my clocks; the sound quality far exceeds typical commercial clock radios. If you have an old pair of speakers gathering dust, why not recycle them and put them to good use with this project?

A switched power output for an

external audio source is enabled when the radio is selected; you could also integrate other audio sources into the enclosure alongside the clock hardware and Raspberry Pi. The switched power output can drive a 5V-powered device directly, feed an external regulator for a lower-voltage device, or act as an open-drain switch to control higher-voltage devices.

Software features

The clock supports many more features than a typical commercial clock radio. Up to twenty alarms can be configured at the same time. The first four alarms can be accessed via the clock's switches, while all alarms can be managed via the web interface.

For each alarm, you set the days of the week, time, duration, media source and relative volume. Any combination of days can be specified, from a single day to all seven. For example, you can set different alarms for weekdays and on the weekend.

Alarms can be configured as one-offs or to repeat indefinitely. One-offs may be defined up to seven days in advance and, if you like, you can set a one-off to occur on all or any of the seven upcoming days. After the one-off trips, it will not recur. I use these when I need to get up early for a work trip.

To confirm you've set your alarm correctly, simultaneously press two alarm selection switches to display the remaining time until the next scheduled alarm. The countdown to the next alarm is displayed for as long as you hold the two alarm selection switches.

An alarm can have a fixed volume, as you'd have with a regular alarm clock, or it can gradually ramp the volume up (or down) in one-minute steps so that you're gently awakened.

When the media or an alarm is playing from a playlist file, the clock will remember the last track so that it continues from the following track next time. While playing, the playlist's contents are visible via the web interface, so you can jump tracks by clicking.

One of the more novel capabilities of this project is its ability to cluster multiple clocks into one system. Clustered clocks share their alarm settings via WiFi, and changes made on one clock will be reflected soon after on the other clustered clocks. Many button press events are also shared in real-time amongst clustered clocks.

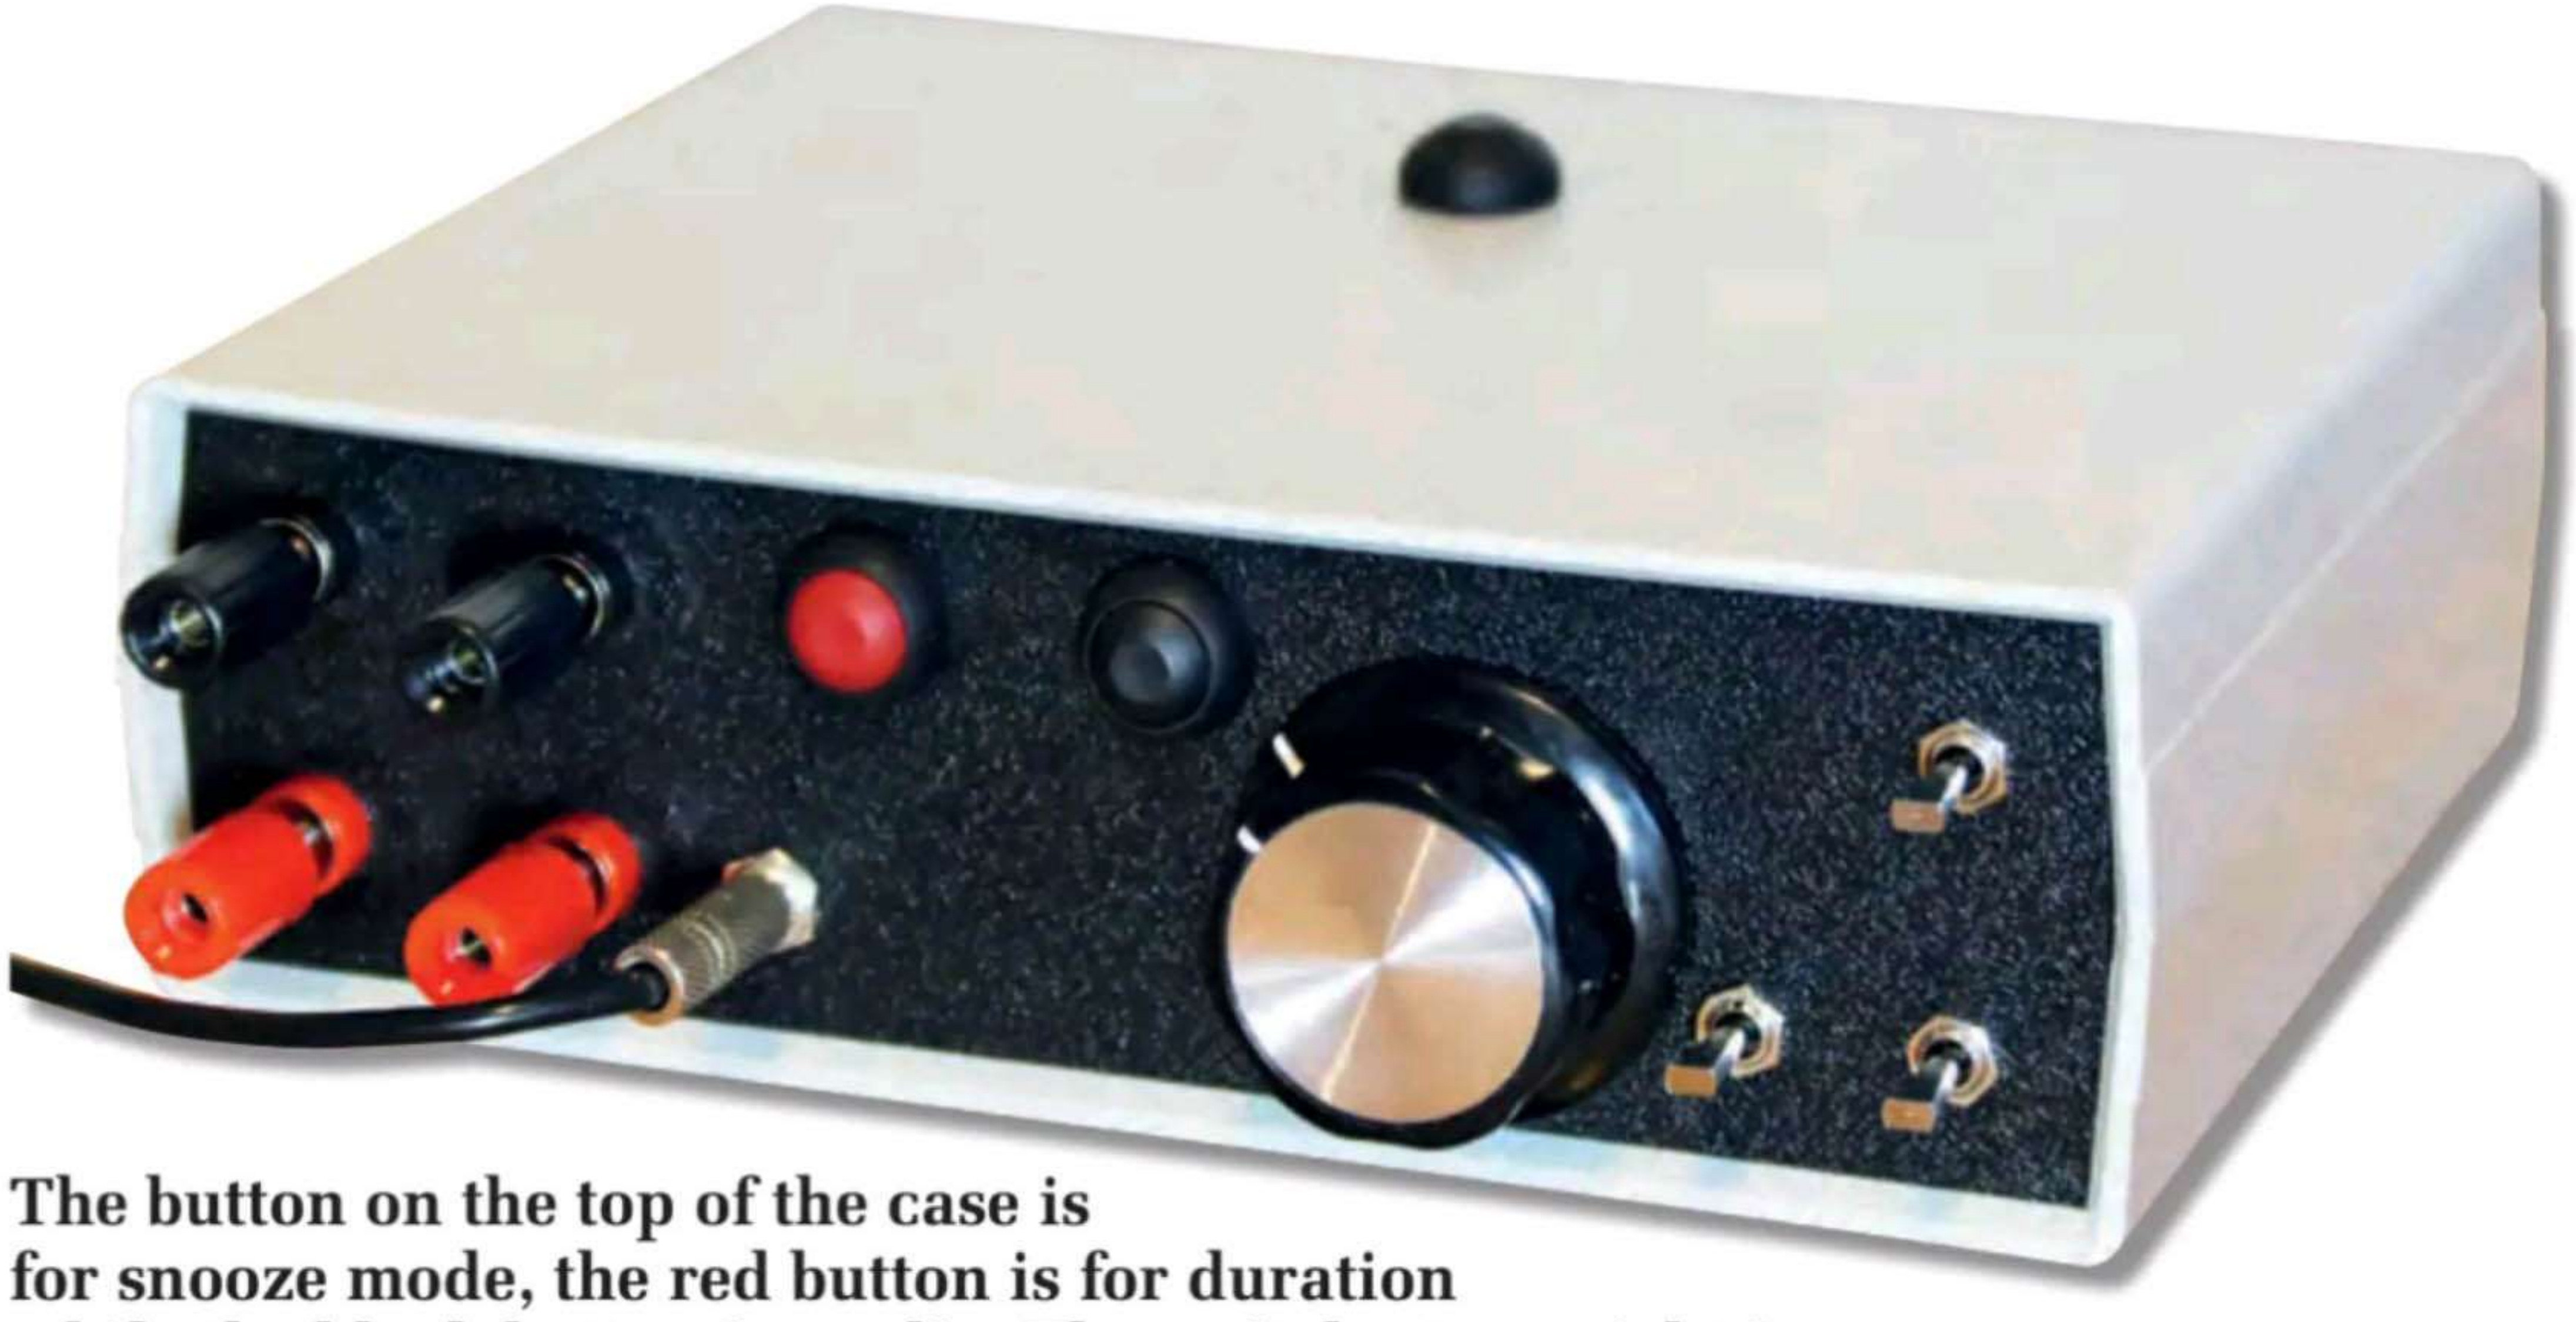

The button on the top of the case is for snooze mode, the red button is for duration while the black button is media. The switch at top right is plus/minus, while the two switches below it handle alarm 1/2 and alarm 3/4 respectively. The big knob in the centre is for radio tuning.

With clocks on either side of the bed, either person can invoke snooze, change volume, modify an alarm etc. You can even use clustering to coordinate clocks in different parts of the home.

Circuit protection features

The circuit includes reverse polarity and overvoltage protection. Raspberry Pis have an absolute maximum supply limit of 6V, beyond which they will be damaged. At our house, we have a box of spare 5V and 12V plugpacks to draw upon for our various devices, and they all share the same style of coaxial connector.

If you accidentally plugged a 12V plugpack into this project, that would instantly destroy it and the Raspberry Pi. The protection circuit was included to guard against that possibility.

Because the circuit mostly follows a 5V design but the Raspberry Pi expansion interface uses strictly 3.3V logic, the clock includes series protection resistances for all general purpose I/O (GPIO) lines to guard against inadvertent shorting to a 5V source. I accidentally did this when prodding around a prototype and was glad for the protection.

Trying the software

You may be interested to try the software, even if you're not yet ready to

WiFi can interfere with Bluetooth

Bluetooth operates in the same 2.4GHz band as WiFi and different devices in that band can interfere with one other, especially when a nearby device is using a lot of bandwidth. Bluetooth interference can cause audio stutter and spontaneous disconnections.

If the interference becomes annoying, reassigning the Bluetooth channel by forgetting all Bluetooth pairings and starting over can help temporarily, as can changing the access point's WiFi channel. However, these strategies may not be effective in the long term.

The Raspberry Pi 4 supports the 5GHz WiFi band, so if Bluetooth audio streaming is an important feature for you, you'll get the best results using a Pi 4 and ensuring there are no 2.4GHz WiFi devices or access points in the same or adjacent rooms as the clock.

You could also disable the 2.4GHz band in all nearby access points and WiFi extenders. However, as most of us have legacy 2.4GHz-only WiFi devices, and sometimes 2.4GHz is the only usable spectrum, implementing this drastic strategy may be difficult (5GHz WiFi doesn't penetrate walls very well).

You can avoid severe interference if your home WiFi is based on recent access point technology supporting both 2.4GHz and 5GHz WiFi bands and band-steering. Configure the band-steering to force 5GHz-capable devices to use 5GHz WiFi channels for the fullest practical signal strength range and check that your 5GHz capable devices have switched over. Also, if possible, use wired Ethernet instead of 2.4GHz backhaul for any WiFi extenders you may have deployed.

Another thing to consider is that microwave ovens operate at around 2.4GHz, so if a kitchen is nearby, an operating microwave oven can interfere with WiFi and Bluetooth in that band.

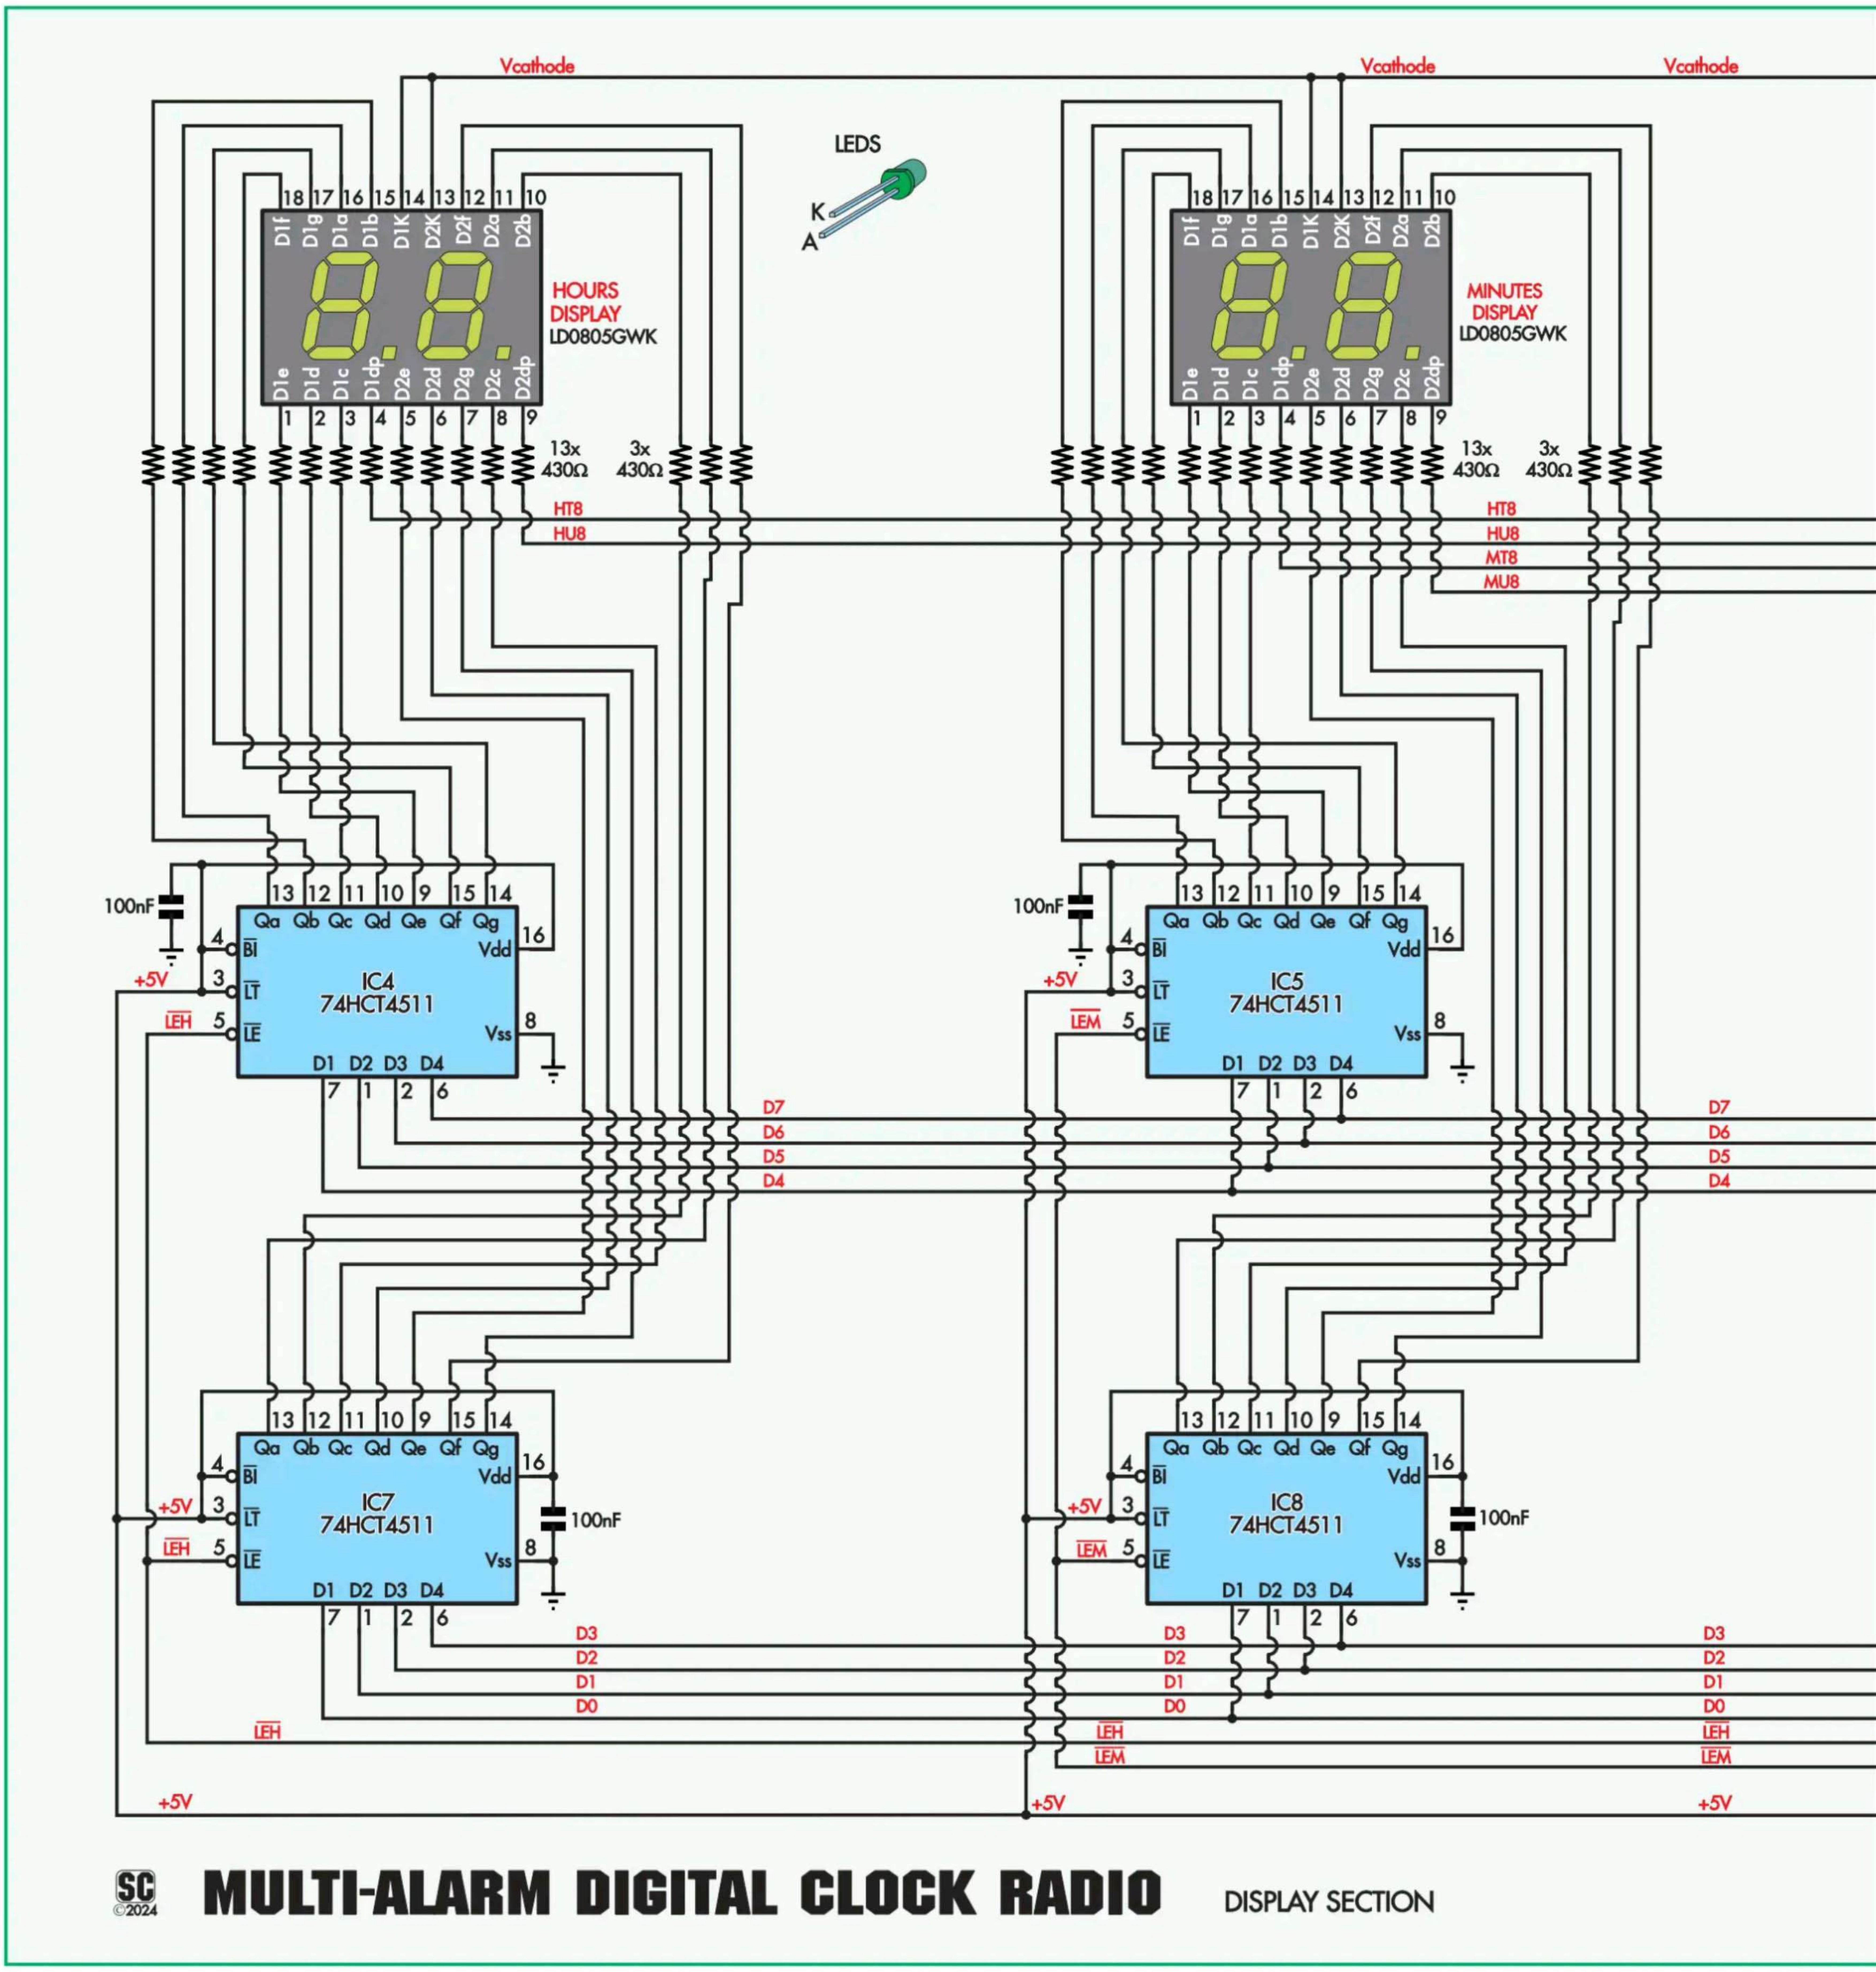
Vcathode
Vcathode
Vcathode
LEDS
K
A
HOURS DISPLAY
LD0805GWK
MINUTES DISPLAY
LD0805GWK
D1f D1g D1a D1b D1K D2K D2f D2a D2b
D1e D1d D1c D1dp D2e D2d D2g D2c D2dp
18 17 16 15 14 13 12 11 10
1 2 3 4 5 6 7 8 9
13x 430Ω
3x 430Ω
HT8
HU8
MT8
MU8
100nF
+5V
LEH
LEM
BI
LT
LE
Qa Qb Qc Qd Qe Qf Qg
Vdd
Vss
D1 D2 D3 D4
IC4 74HCT4511
IC5 74HCT4511
IC7 74HCT4511
IC8 74HCT4511
D7
D6
D5
D4
D3
D2
D1
D0
SC ©2024
MULTI-ALARM DIGITAL CLOCK RADIO
DISPLAY SECTION

commit to the construction. The software can be installed onto any Linux-based Raspberry Pi with a 40 pin expansion connector (the GPIO library currently does not support the Pi 5).

Without the clock hardware, you can use the web GUI to set up and configure alarms, watch the alarms trip, pair your phone or tablet with the 'clock' to use it as a Bluetooth speaker or play media from the Pi's flash card, an attached USB drive, a network share or from the internet.

A script simplifies installing and configuring the Pi. It fetches the required libraries, installs them, then configures the clock, a file server, web interface, media player, automatic updates, NTP and time monitoring processes.

You can optionally enable a firewall so the clock cannot be accessed from outside your home network and/or turn off the Pi's power and activity LEDs so they don't keep you awake at night.

You can download the ZIP file from **siliconchip.au/Shop/6/278** containing a Linux 'tarball' of the software and a PDF document explaining how to prepare the SD card, copy the tarball and

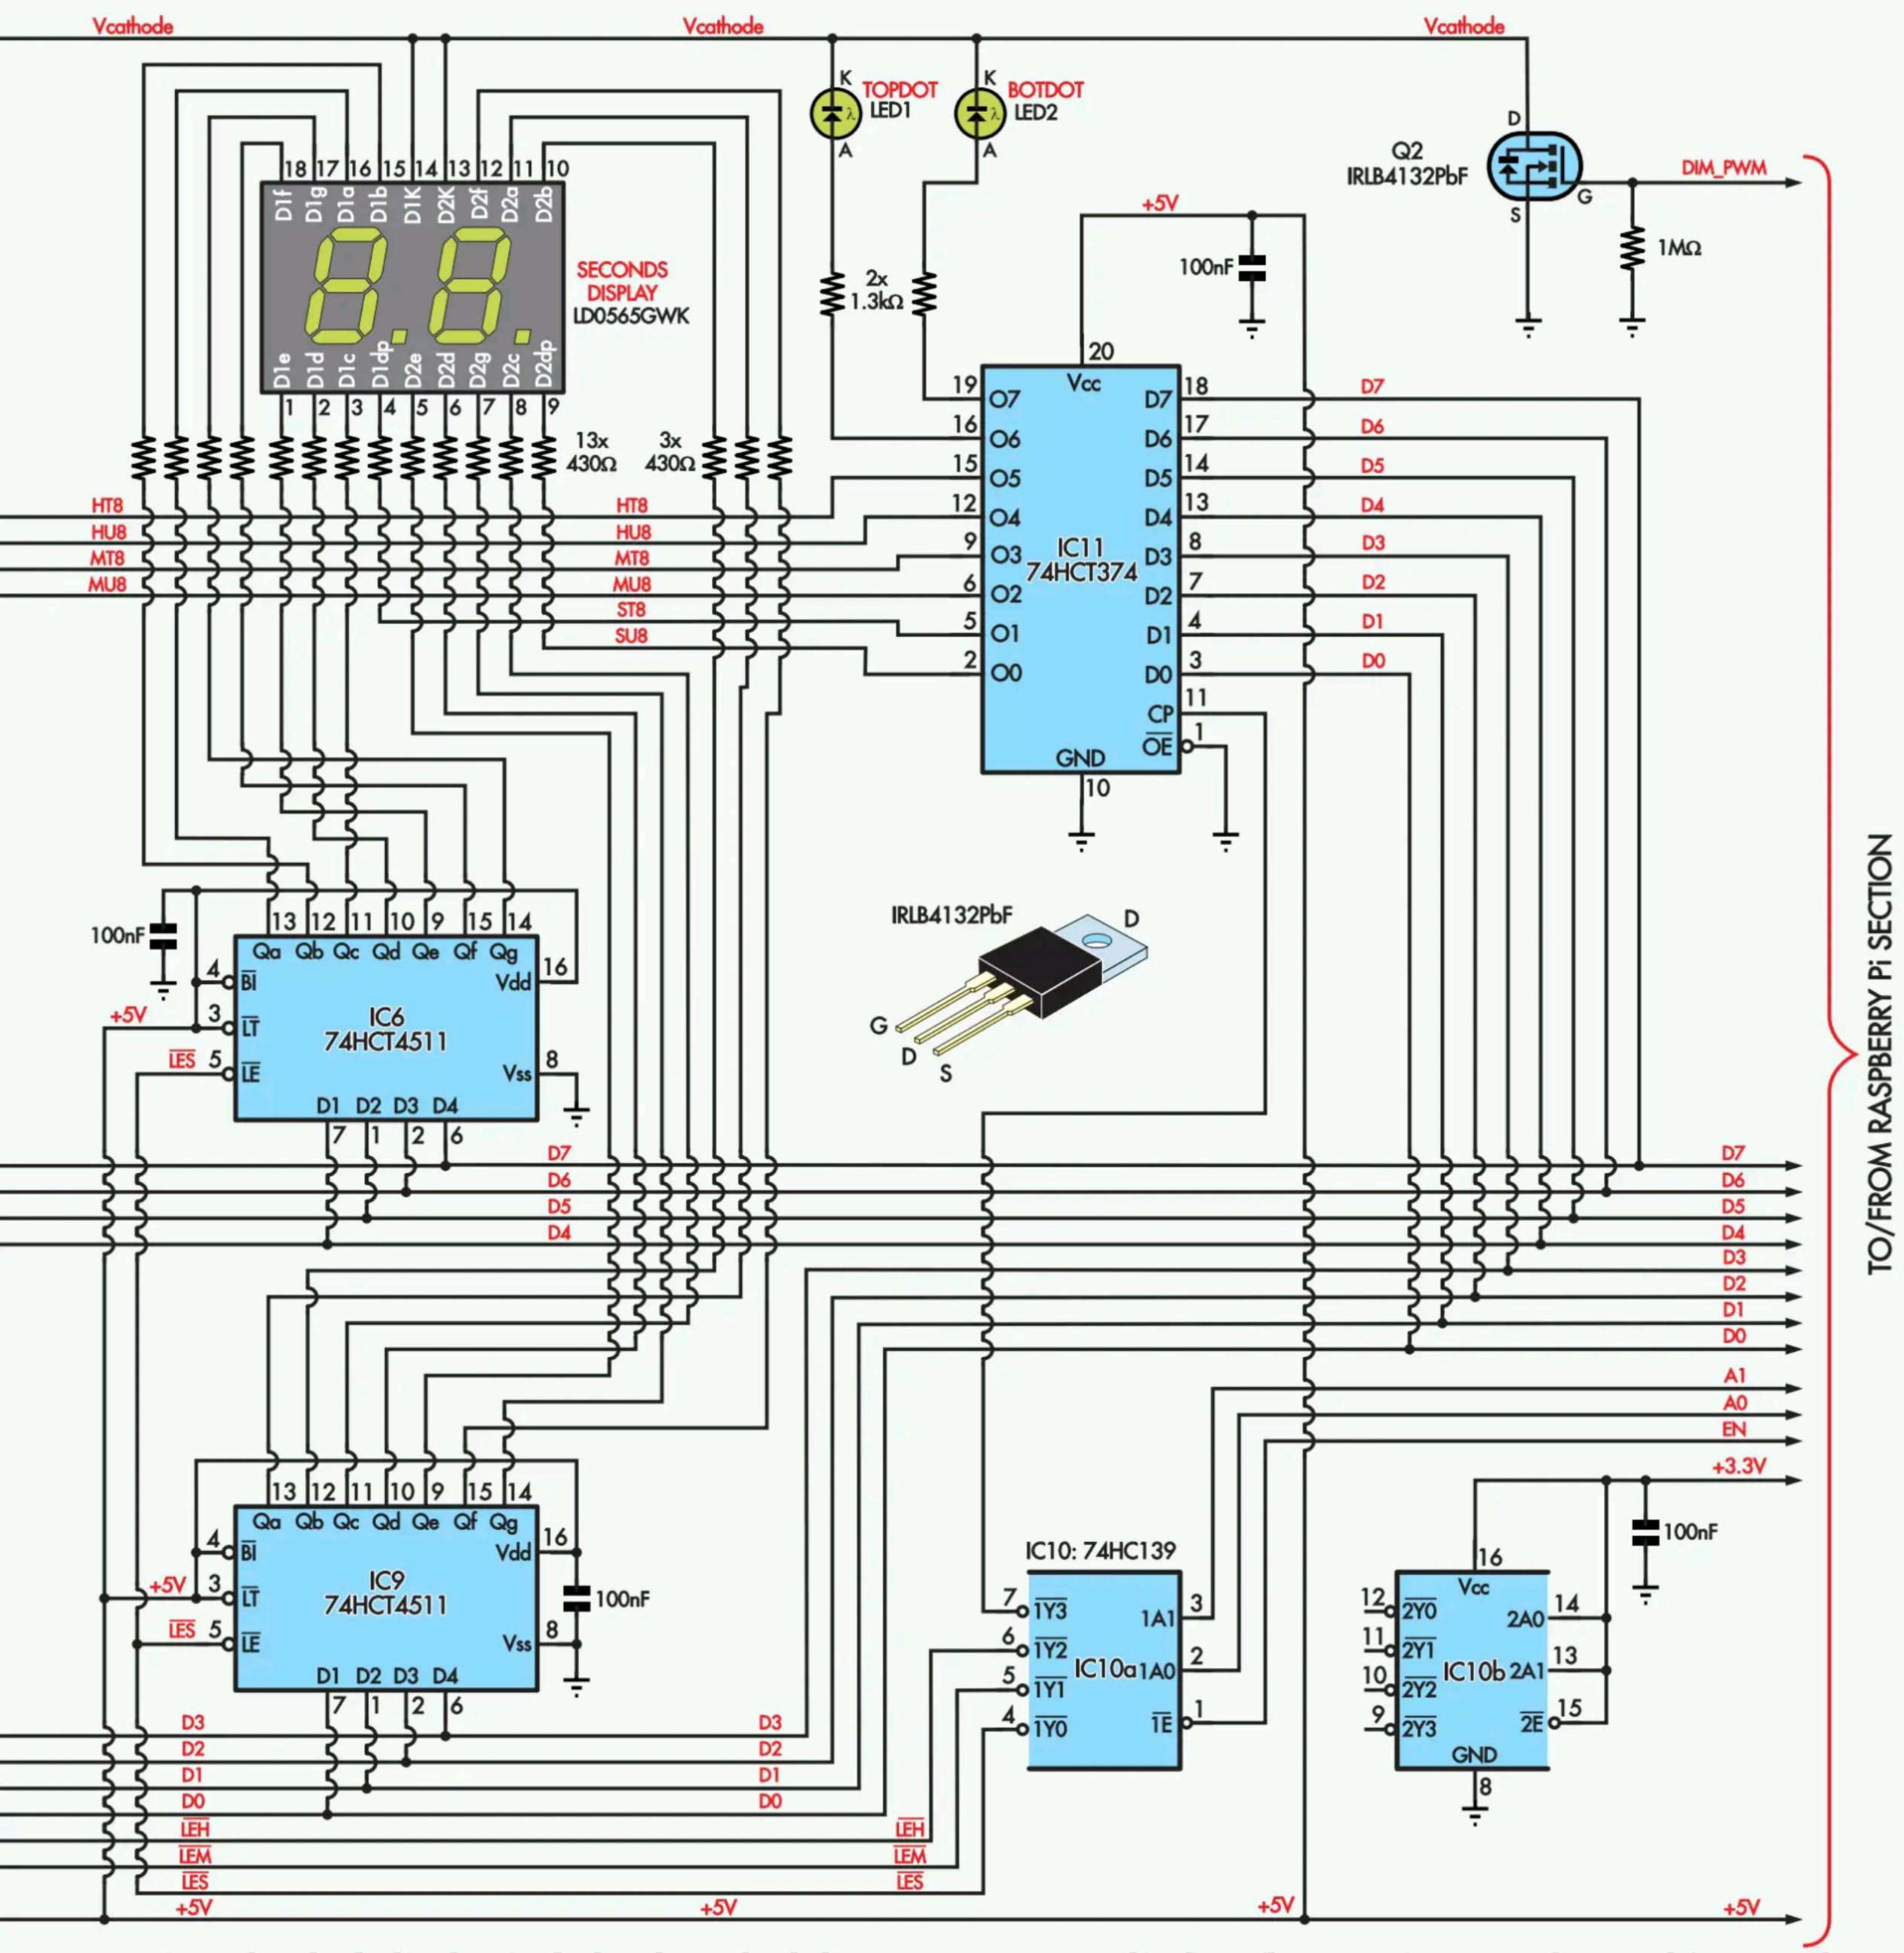
Vcathode
SECONDS DISPLAY LD0565GWK
D1f D1g D1a D1b D1K D2K D2f D2a D2b
D1e D1d D1c D1dp D2e D2d D2g D2c D2dp
13x 430Ω
3x 430Ω
TOPDOT LED1
BOTDOT LED2
K
A
2x 1.3kΩ
Q2 IRLB4132PbF
D
G
S
DIM_PWM
1MΩ
+5V
100nF
IC11 74HCT374
Vcc
O7 O6 O5 O4 O3 O2 O1 O0
D7 D6 D5 D4 D3 D2 D1 D0
CP
OE
GND
HT8
HU8
MT8
MU8
ST8
SU8
IRLB4132PbF
IC6 74HCT4511
IC9 74HCT4511
Qa Qb Qc Qd Qe Qf Qg
BI
LT
LE
Vdd
Vss
D1 D2 D3 D4
LES
LEH
LEM
IC10: 74HC139
IC10a
1Y3 1Y2 1Y1 1Y0
1A1 1A0 1E
IC10b
2Y0 2Y1 2Y2 2Y3
2A0 2A1 2E
GND
A1
A0
EN
+3.3V
TO/FROM RASPBERRY Pi SECTION

Fig.1: the clock display includes three dual-digit seven-segment displays (hours, minutes and seconds), two colon LEDs, eight ICs to drive the LEDs and Mosfet Q2 for PWM display brightness control. IC4-IC9 are seven-segment display drivers, while IC11 is an eight-bit latch that drives the decimal points and colons.

run the installation script. There are also notes about software debugging modes for testing.

See the panel on page 16 of this article for instructions on installing the software.

Circuit details

The Clock Radio circuit diagram is shown in Figs.1 & 2. Fig.1 is the display section with the LED arrays and their drivers. That section is driven by the control section shown in Fig.2, which also has the audio, user interface (switch/button) and power portions.

The 5V and 3.3V power rails for the display circuitry shown in Fig.1 come from the Raspberry Pi controller in Fig.2, along with the following digital data lines via 1kΩ resistors: an 8-bit data bus (D0-D7), a two-bit address bus (A0 & A1), a latch signal (EN) and a PWM brightness control line (DIM_PWM).

By setting the eight data lines and the address, then 'strobing' (pulsing) the latch, the software on the Pi can update the digits for the hours, minutes and seconds, the six decimal points and two colon LEDs.

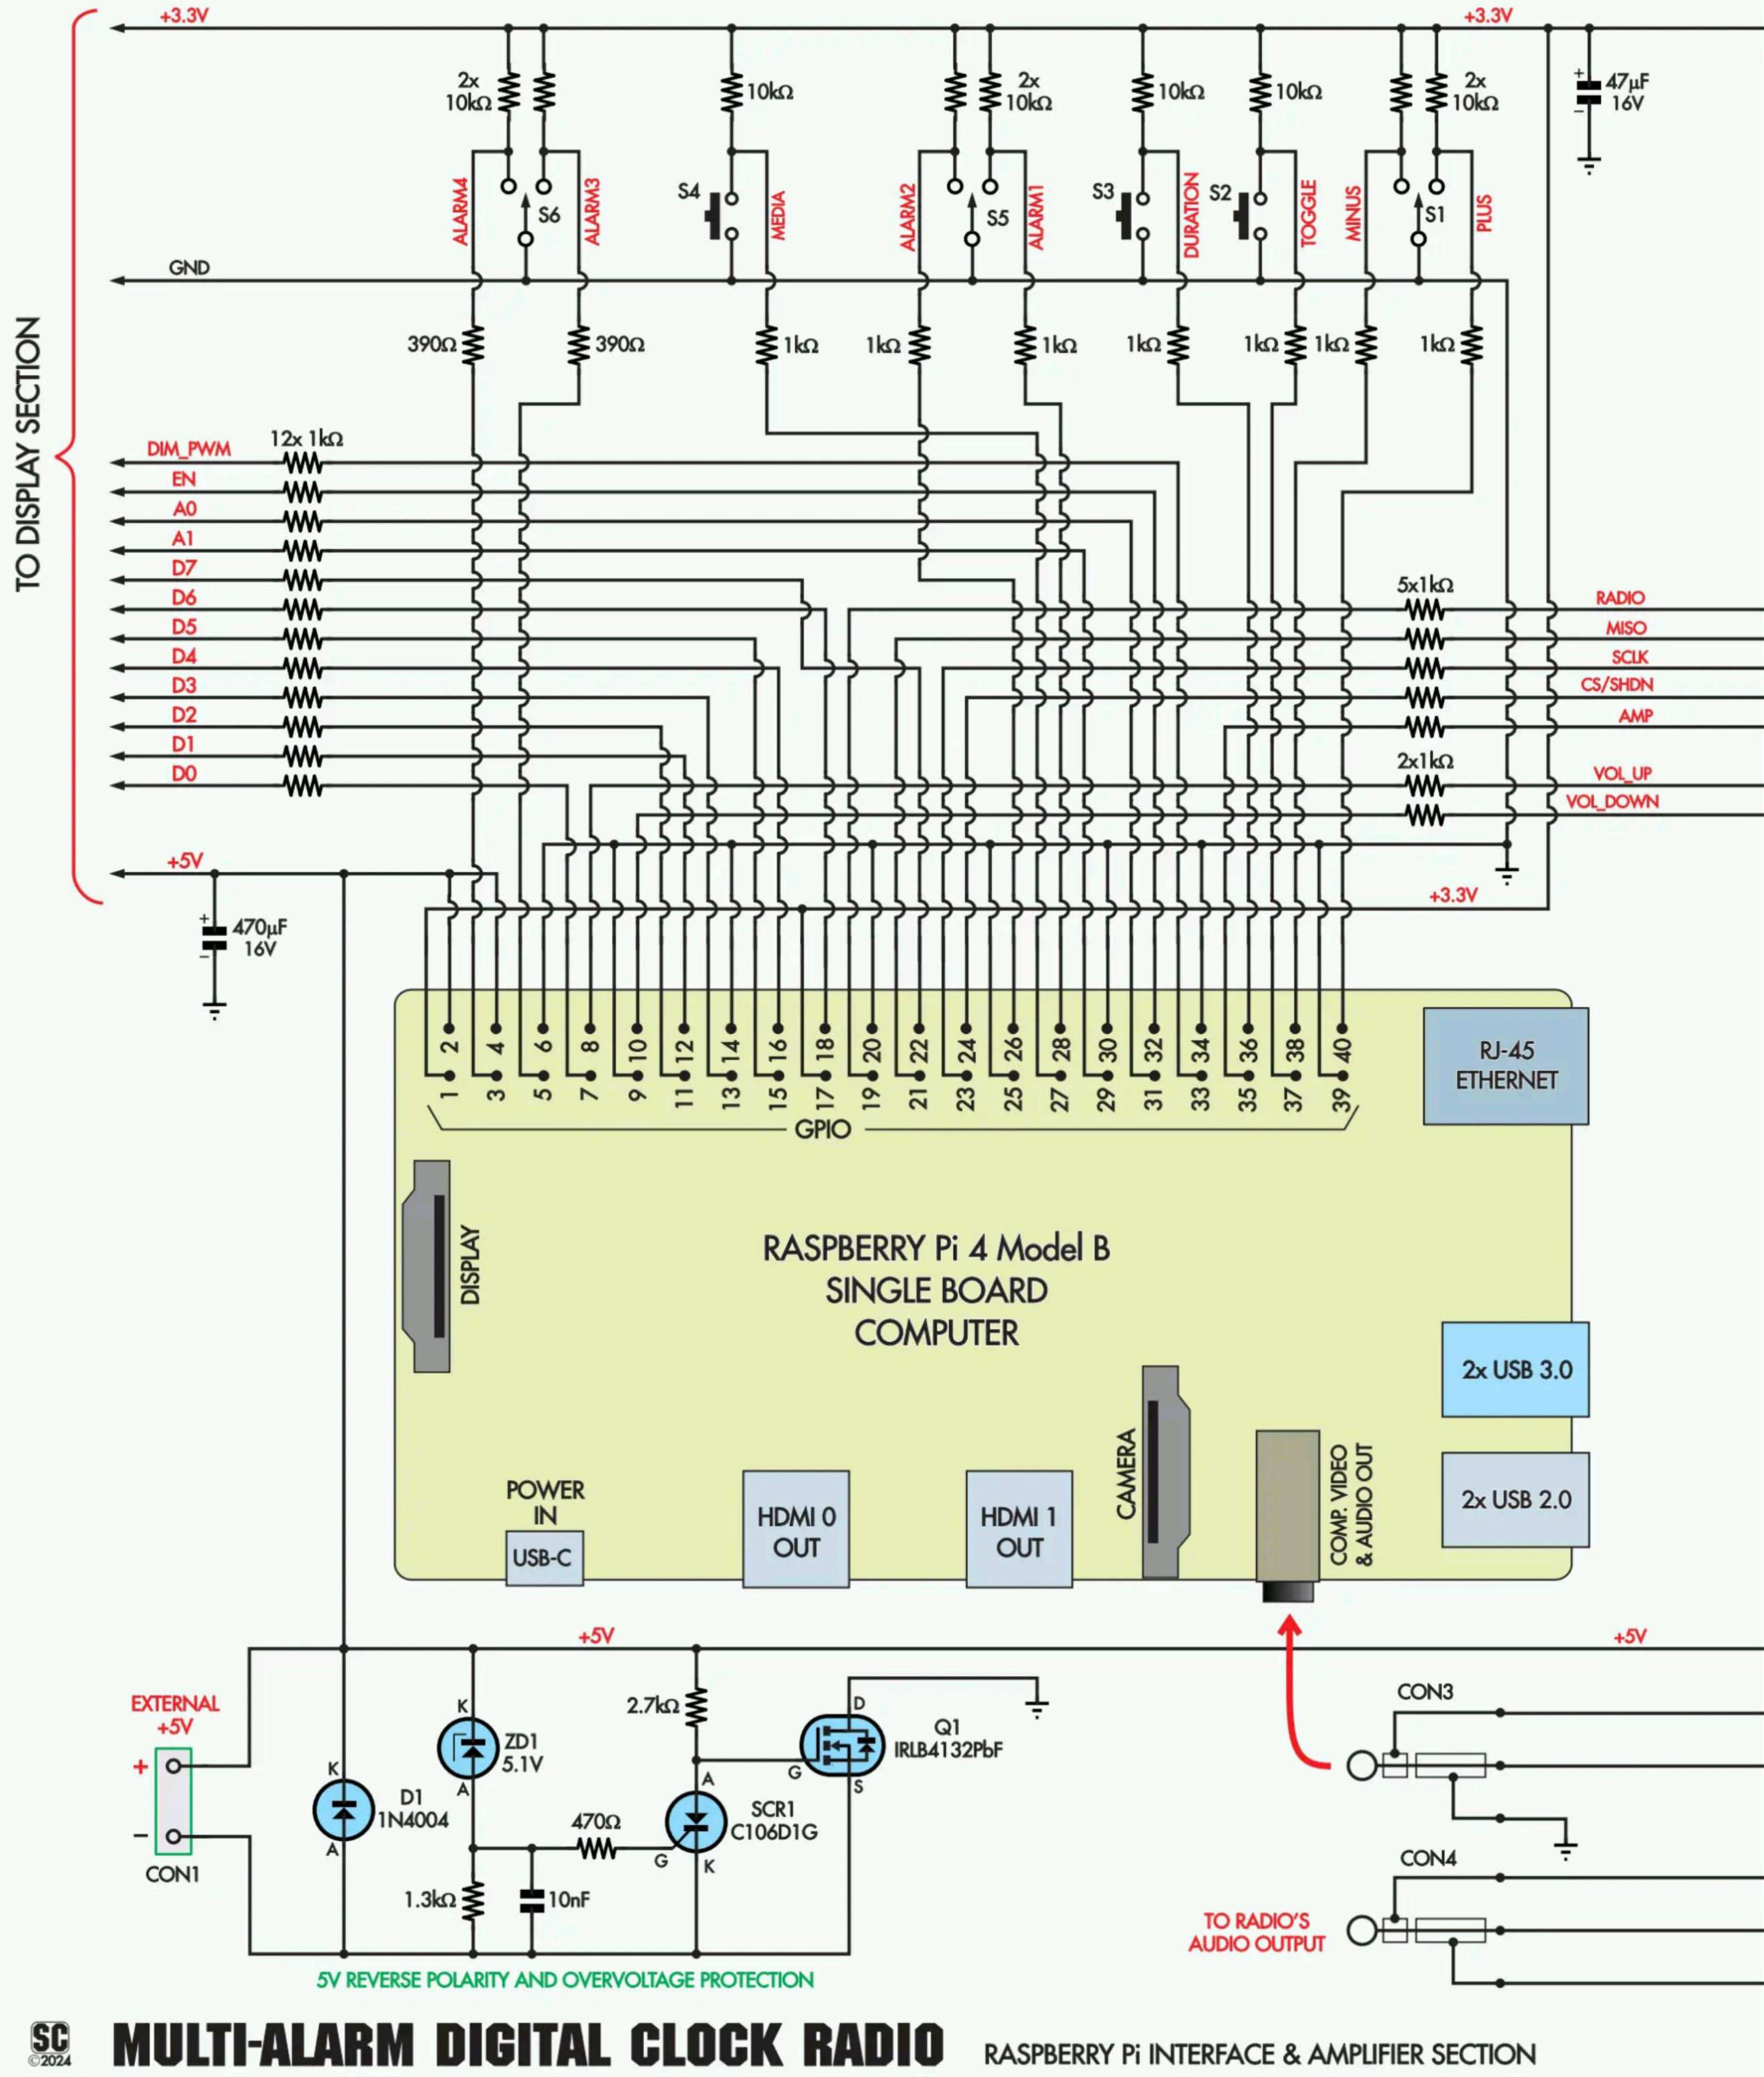
+3.3V
GND
+5V
TO DISPLAY SECTION
2x 10kΩ
10kΩ
ALARM4
S6
ALARM3
S4
MEDIA
ALARM2
S5
ALARM1
S3
DURATION
S2
TOGGLE
MINUS
S1
PLUS
390Ω
1kΩ
47µF 16V
12x 1kΩ
DIM_PWM
EN
A0
A1
D7
D6
D5
D4
D3
D2
D1
D0
5x1kΩ
RADIO
MISO
SCLK
CS/SHDN
AMP
2x1kΩ
VOL_UP
VOL_DOWN
470µF 16V
GPIO
RJ-45 ETHERNET
DISPLAY
RASPBERRY Pi 4 Model B SINGLE BOARD COMPUTER
2x USB 3.0
2x USB 2.0
CAMERA
COMP. VIDEO & AUDIO OUT
POWER IN
USB-C
HDMI 0 OUT
HDMI 1 OUT
EXTERNAL +5V
CON1
D1 1N4004
ZD1 5.1V
2.7kΩ
470Ω
1.3kΩ
10nF
SCR1 C106D1G
Q1 IRLB4132PbF
CON3
CON4
TO RADIO'S AUDIO OUTPUT
5V REVERSE POLARITY AND OVERVOLTAGE PROTECTION
SC ©2024
MULTI-ALARM DIGITAL CLOCK RADIO
RASPBERRY Pi INTERFACE & AMPLIFIER SECTION

Fig.2: the Raspberry Pi connects to the display circuitry shown in Fig.1 using 12 digital lines that go via 1kΩ resistors. The switches and buttons also connect to the Pi's digital I/O pins with pull-up resistors, while the ambient brightness monitoring and audio amplification circuitry are at upper right. The section at bottom left protects against power supply over-voltage and reversed polarity.

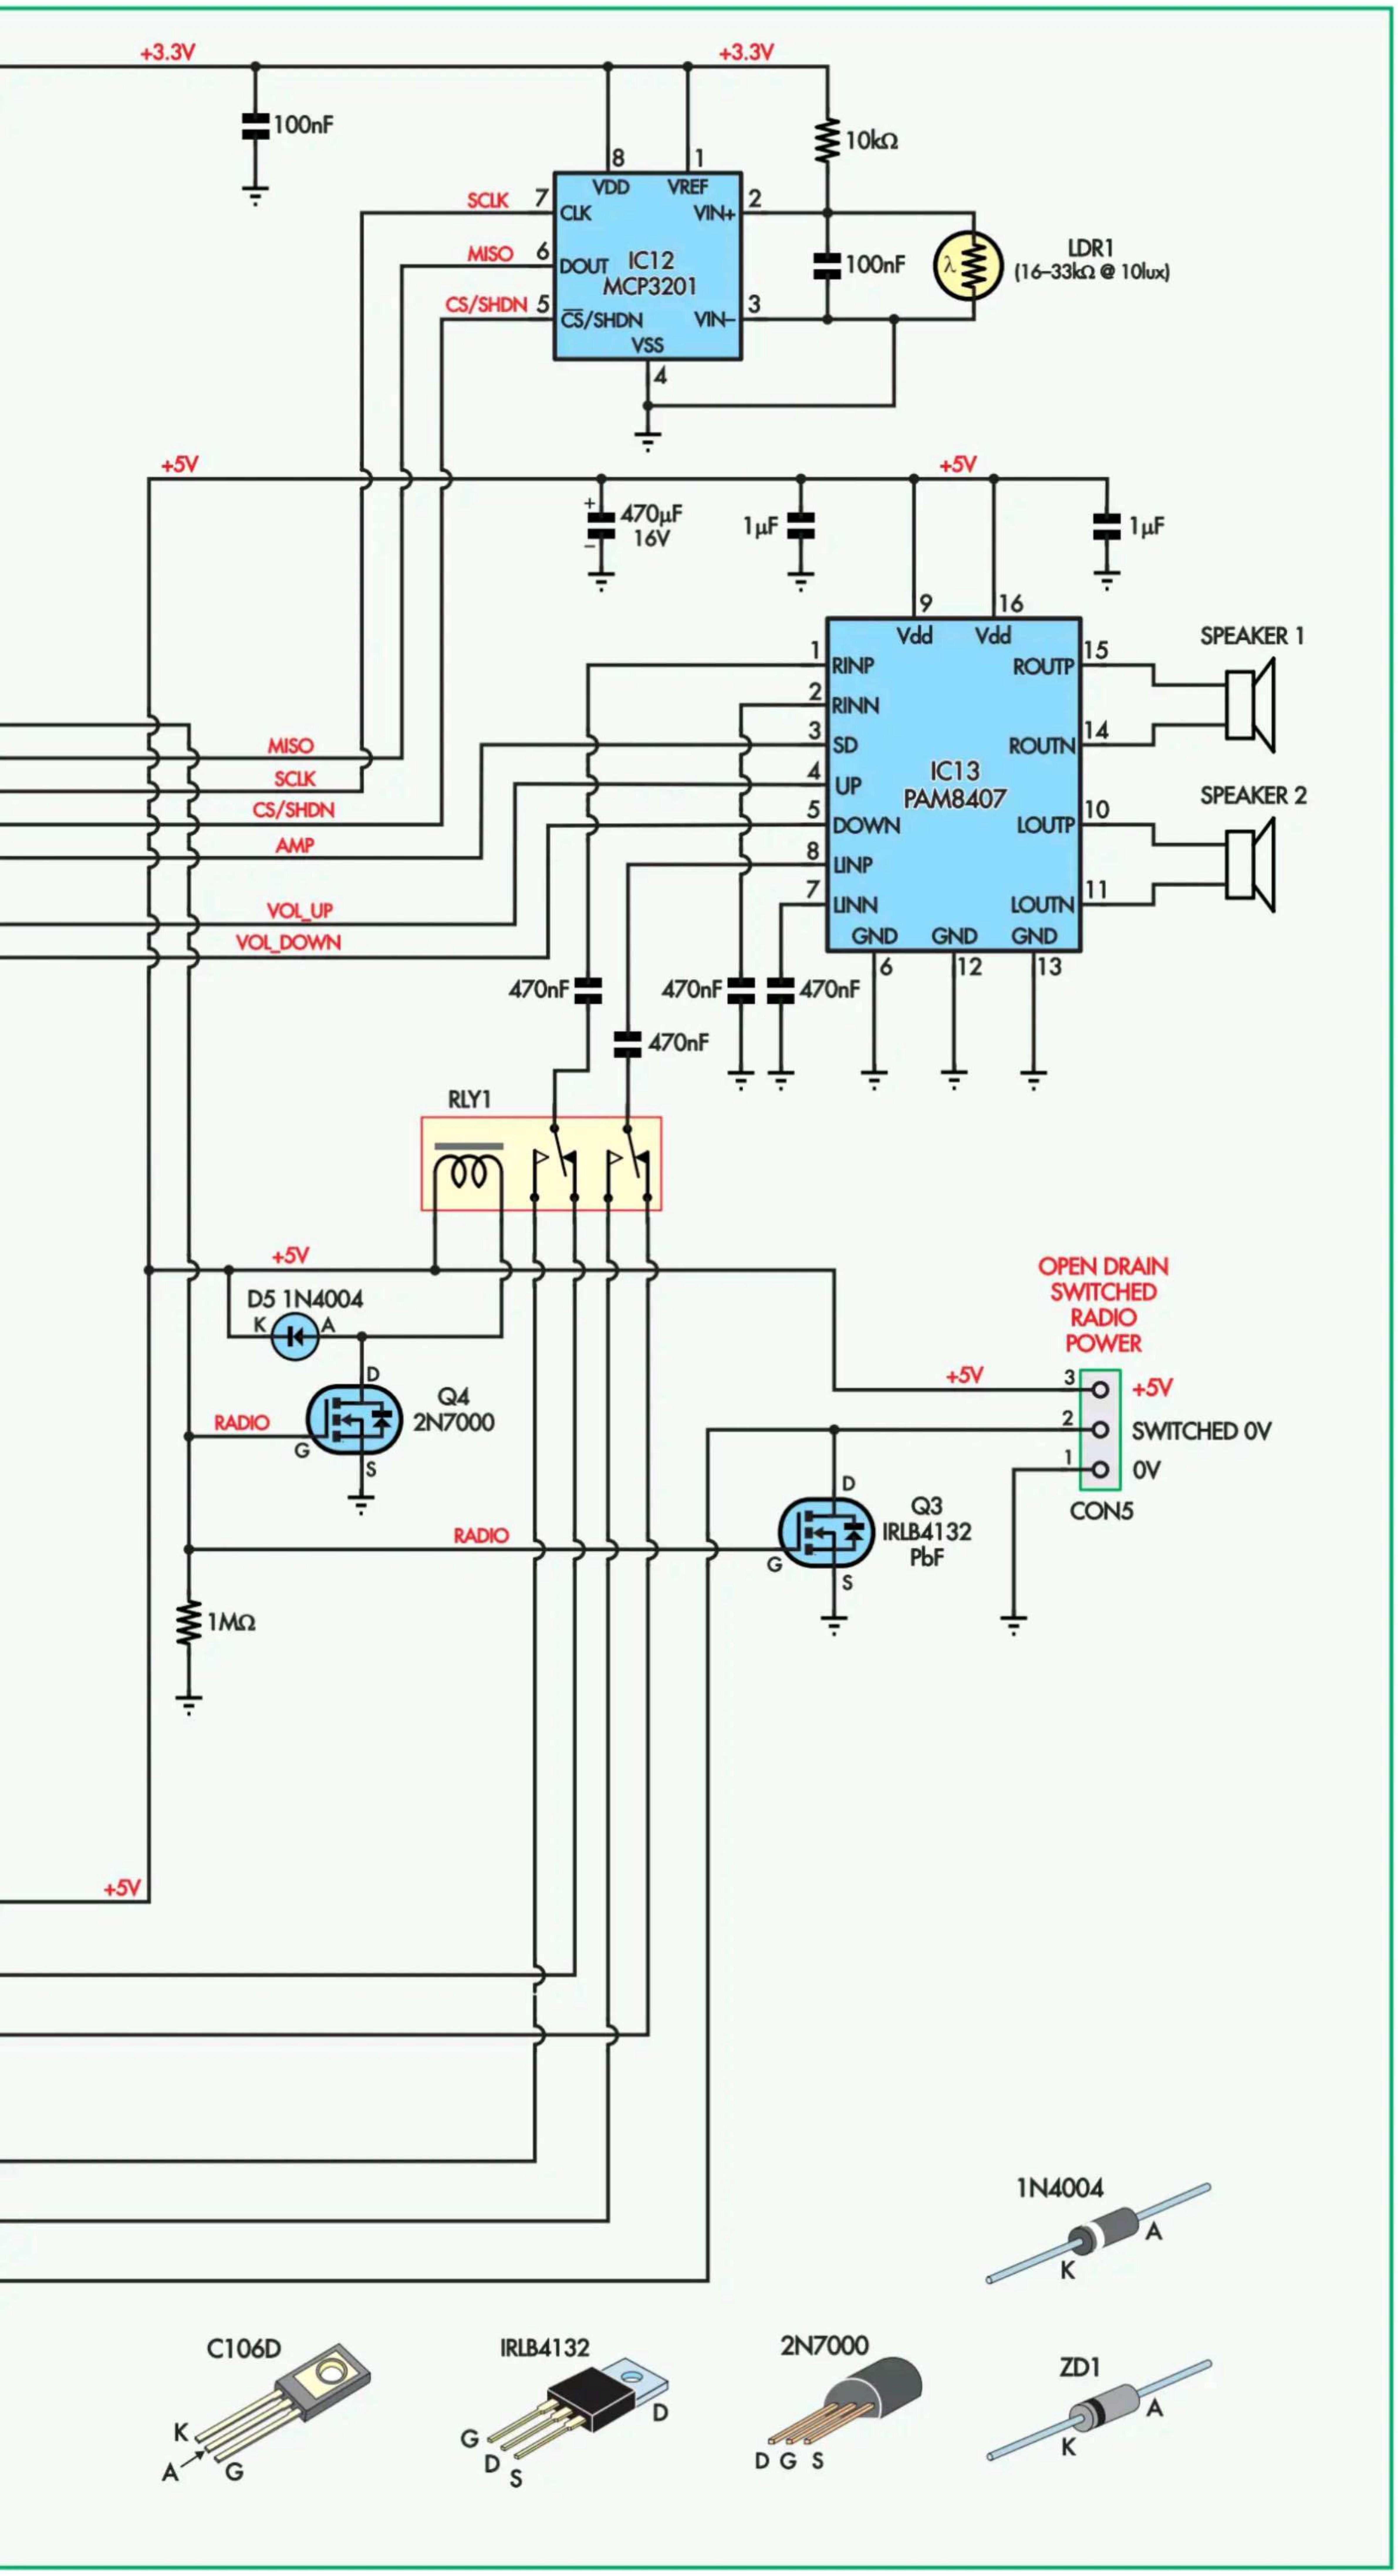
+3.3V
100nF
+3.3V
10kΩ
SCLK
MISO
CS/SHDN
IC12
MCP3201
VDD
VREF
CLK
VIN+
DOUT
CS/SHDN
VIN–
VSS
100nF
LDR1
(16–33kΩ @ 10lux)
+5V
470μF
16V
1μF
+5V
1μF
Vdd
Vdd
RINP
RINN
SD
UP
DOWN
LINP
LINN
GND
GND
GND
ROUTP
ROUTN
LOUTP
LOUTN
IC13
PAM8407
SPEAKER 1
SPEAKER 2
MISO
SCLK
CS/SHDN
AMP
VOL_UP
VOL_DOWN
470nF
470nF
470nF
470nF
RLY1
+5V
D5 1N4004
K
A
D
G
S
Q4
2N7000
RADIO
OPEN DRAIN SWITCHED RADIO POWER
+5V
+5V
SWITCHED 0V
0V
CON5
D
Q3
IRLB4132
PbF
G
S
RADIO
1MΩ
+5V
1N4004
A
K
ZD1
A
K
C106D
K
A
G
IRLB4132
G
D
S
D
2N7000
D G S

The seven-segment displays are driven by six BCD-to-seven-segment display drivers, IC4 to IC9, and the dots and colon from IC11. IC4 to IC9 convert binary numbers to segment patterns on the seven segment displays and can deliver the necessary LED drive current. IC11 works like a one byte (eight bit) memory to remember which dots are turned on.

These chips have 3.3V-compatible inputs, suiting the Pi's GPIO bus, and 5V outputs that can draw from the higher-current 5V supply rail. It is important to use 74HCT chips rather than 74HC because the latter are marginal at recognising 3.3V as a high level while the former have a maximum high threshold of 2V.

Decoding the address bus and latching of the data is performed by IC10. As the decoding logic is all at the same level (3.3V), IC10 can be of the 74HC variety.

IC4 to IC9 and IC11 drive all the LED display anodes via nominally 430Ω resistors while the LED display cathodes all go to the drain of N-channel Mosfet Q2.

A PWM signal applied to Q2's gate therefore determines the overall brightness of all the LEDs. A 1MΩ resistor holds it off whenever the Pi is not actively driving it, so the display is blank when the Pi software is not running.

The software cannot determine whether all LEDs are present because the display section is a 'write-only interface'. If you don't need them, you could leave off the seconds LEDs and associated BCD driver chips, and no software changes will be required.

Matching LED brightness

Theoretically, identical displays from the same vendor should have the same brightness. As the project uses a combination of 0.8-inch 7-segment displays, 0.56-inch 7-segment displays and discrete LEDs, they might not all be the same efficiency. In that case, they can be equalised by adjusting the values of the 430Ω current-limiting resistors.

Four of the five clock prototypes used Lumex 7-segment displays, and both sizes produced identical brightnesses. One prototype used Multicomp Pro devices, resulting in the smaller digits being slightly brighter than the larger digits. The larger Multicomp Pro displays were slightly less bright than the equivalent Lumex devices, but the clock's brightness adjustment had the headroom to compensate.

To equalise the Multicomp Pro display intensities, I changed the smaller display's current limiting resistors to 820Ω on that Clock Radio.

If you construct the board using Multicomp Pro parts, we suggest not populating the small display's current limiting resistors until you've built and tested your clock and can determine the optimal resistance. If constructing with Lumex, as Dirty Harry said, you've got to ask yourself a question: "Do I feel lucky?".

You can populate the 430Ω resistors for the small display as we did, but there's a chance you might need to adjust them if they don't match adequately (we didn't need to).

The two discrete 3mm LEDs that make up the colon (":") have characteristics independent from the 7-segment displays. For the devices specified in the parts list, we found 1.3kΩ series resistors illuminated the colon about the same as the 7-segment displays from either vendor. Any 3mm LEDs will work in this design, but be prepared to experiment with those resistor values if you use different parts.

Dimming

The dimming function of the circuit comprises an ambient light level monitor and the PWM control mentioned above. The ambient light level is sampled by a light-dependent resistor (LDR), which forms a voltage divider with a 10kΩ resistor across the 3.3V rail. The brighter the ambient light level, the lower the LDR/resistor junction voltage.

IC12 is an MCP3201 12-bit analog-to-digital converter (ADC) used to measure this voltage. The raw number read from the ADC becomes smaller as the ambient light level increases; the software processes it into a value with 0 indicating darkness and 4095 being the maximum measurable brightness, as shown on the web setup page.

The MCP3201 comes in two versions with different accuracies labelled B & C. You can save yourself a dollar because the cheaper, less-accurate part (C) works fine in this circuit.

The parts list specifies two LDRs that will work well. Ideally, the LDR dark resistance should be at least 10 times its light resistance. The setup page on the web GUI includes four sliders for adjusting the minimum and maximum LED brightness and specifying the corresponding LDR levels.

The sliders provide a lot of flexibility to adjust for minor differences in LDR characteristics so that the display achieves the full range of potential LED brightness. If you choose a different LDR and can't get the dimming to work over the whole range, the 10kΩ resistor value will need to change.

In response to the ambient light level, the software generates a 50Hz PWM waveform that drives the gate of Mosfet Q2 and continuously updates the PWM duty cycle according to the ambient light measurements.

Although the Pi has two high-resolution timers that could be used for hardware PWM timing, neither is available in this design. One is used for the Pi's analog audio output, while the other is commandeered by the Pi's GPIO daemon (service).

The LED brightness PWM is therefore generated in software by the GPIO daemon. You're unlikely to notice that; the worst case is when the display is at its dimmest and the CPU is heavily loaded, such as when an alarm has tripped and it is downloading, decompressing and playing a media file. In that case, the software reduces the PWM frequency to minimise the jitter induced in the PWM signal.

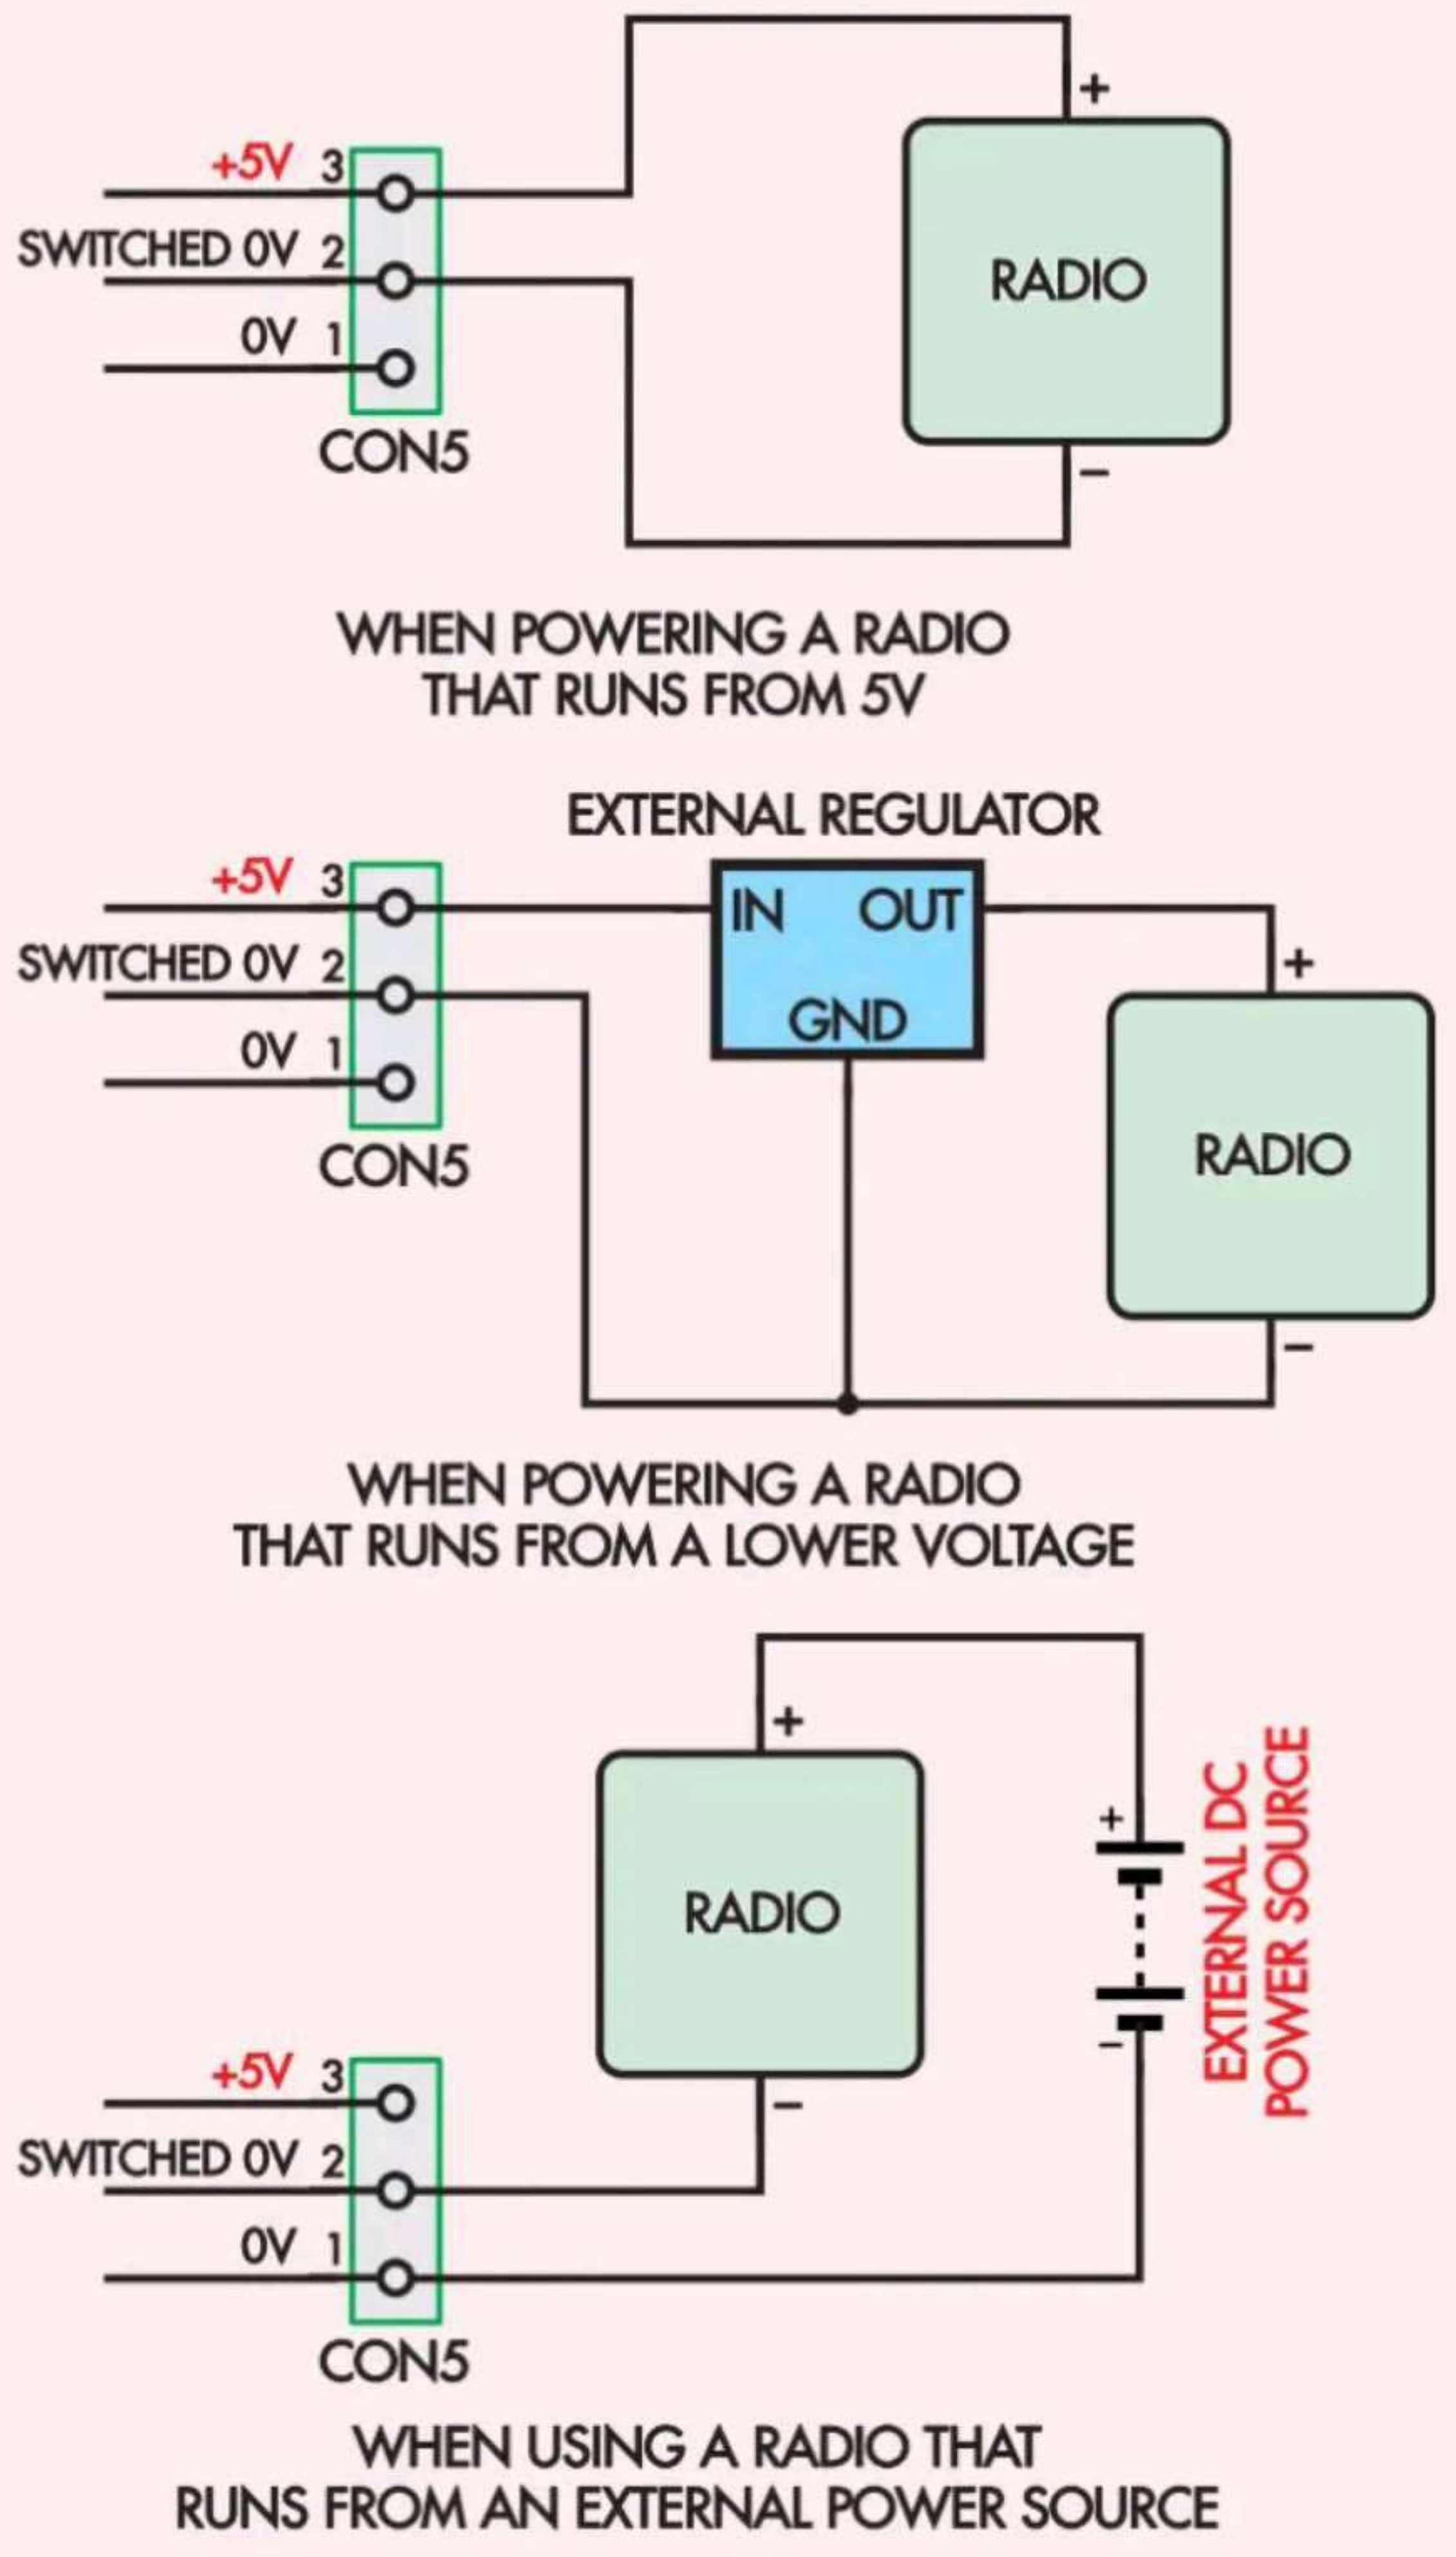
+5V 3
SWITCHED 0V 2
0V 1
CON5
+
RADIO
–
WHEN POWERING A RADIO THAT RUNS FROM 5V
EXTERNAL REGULATOR
+5V 3
SWITCHED 0V 2
0V 1
IN
OUT
GND
+
RADIO
–
CON5
WHEN POWERING A RADIO THAT RUNS FROM A LOWER VOLTAGE
+
RADIO
–
+
EXTERNAL DC POWER SOURCE
+5V 3
SWITCHED 0V 2
0V 1
CON5
WHEN USING A RADIO THAT RUNS FROM AN EXTERNAL POWER SOURCE

Fig.3: you can add a radio receiver board, which will only be powered on when needed, via CON5. Here are three ways to connect it depending on its power requirements.

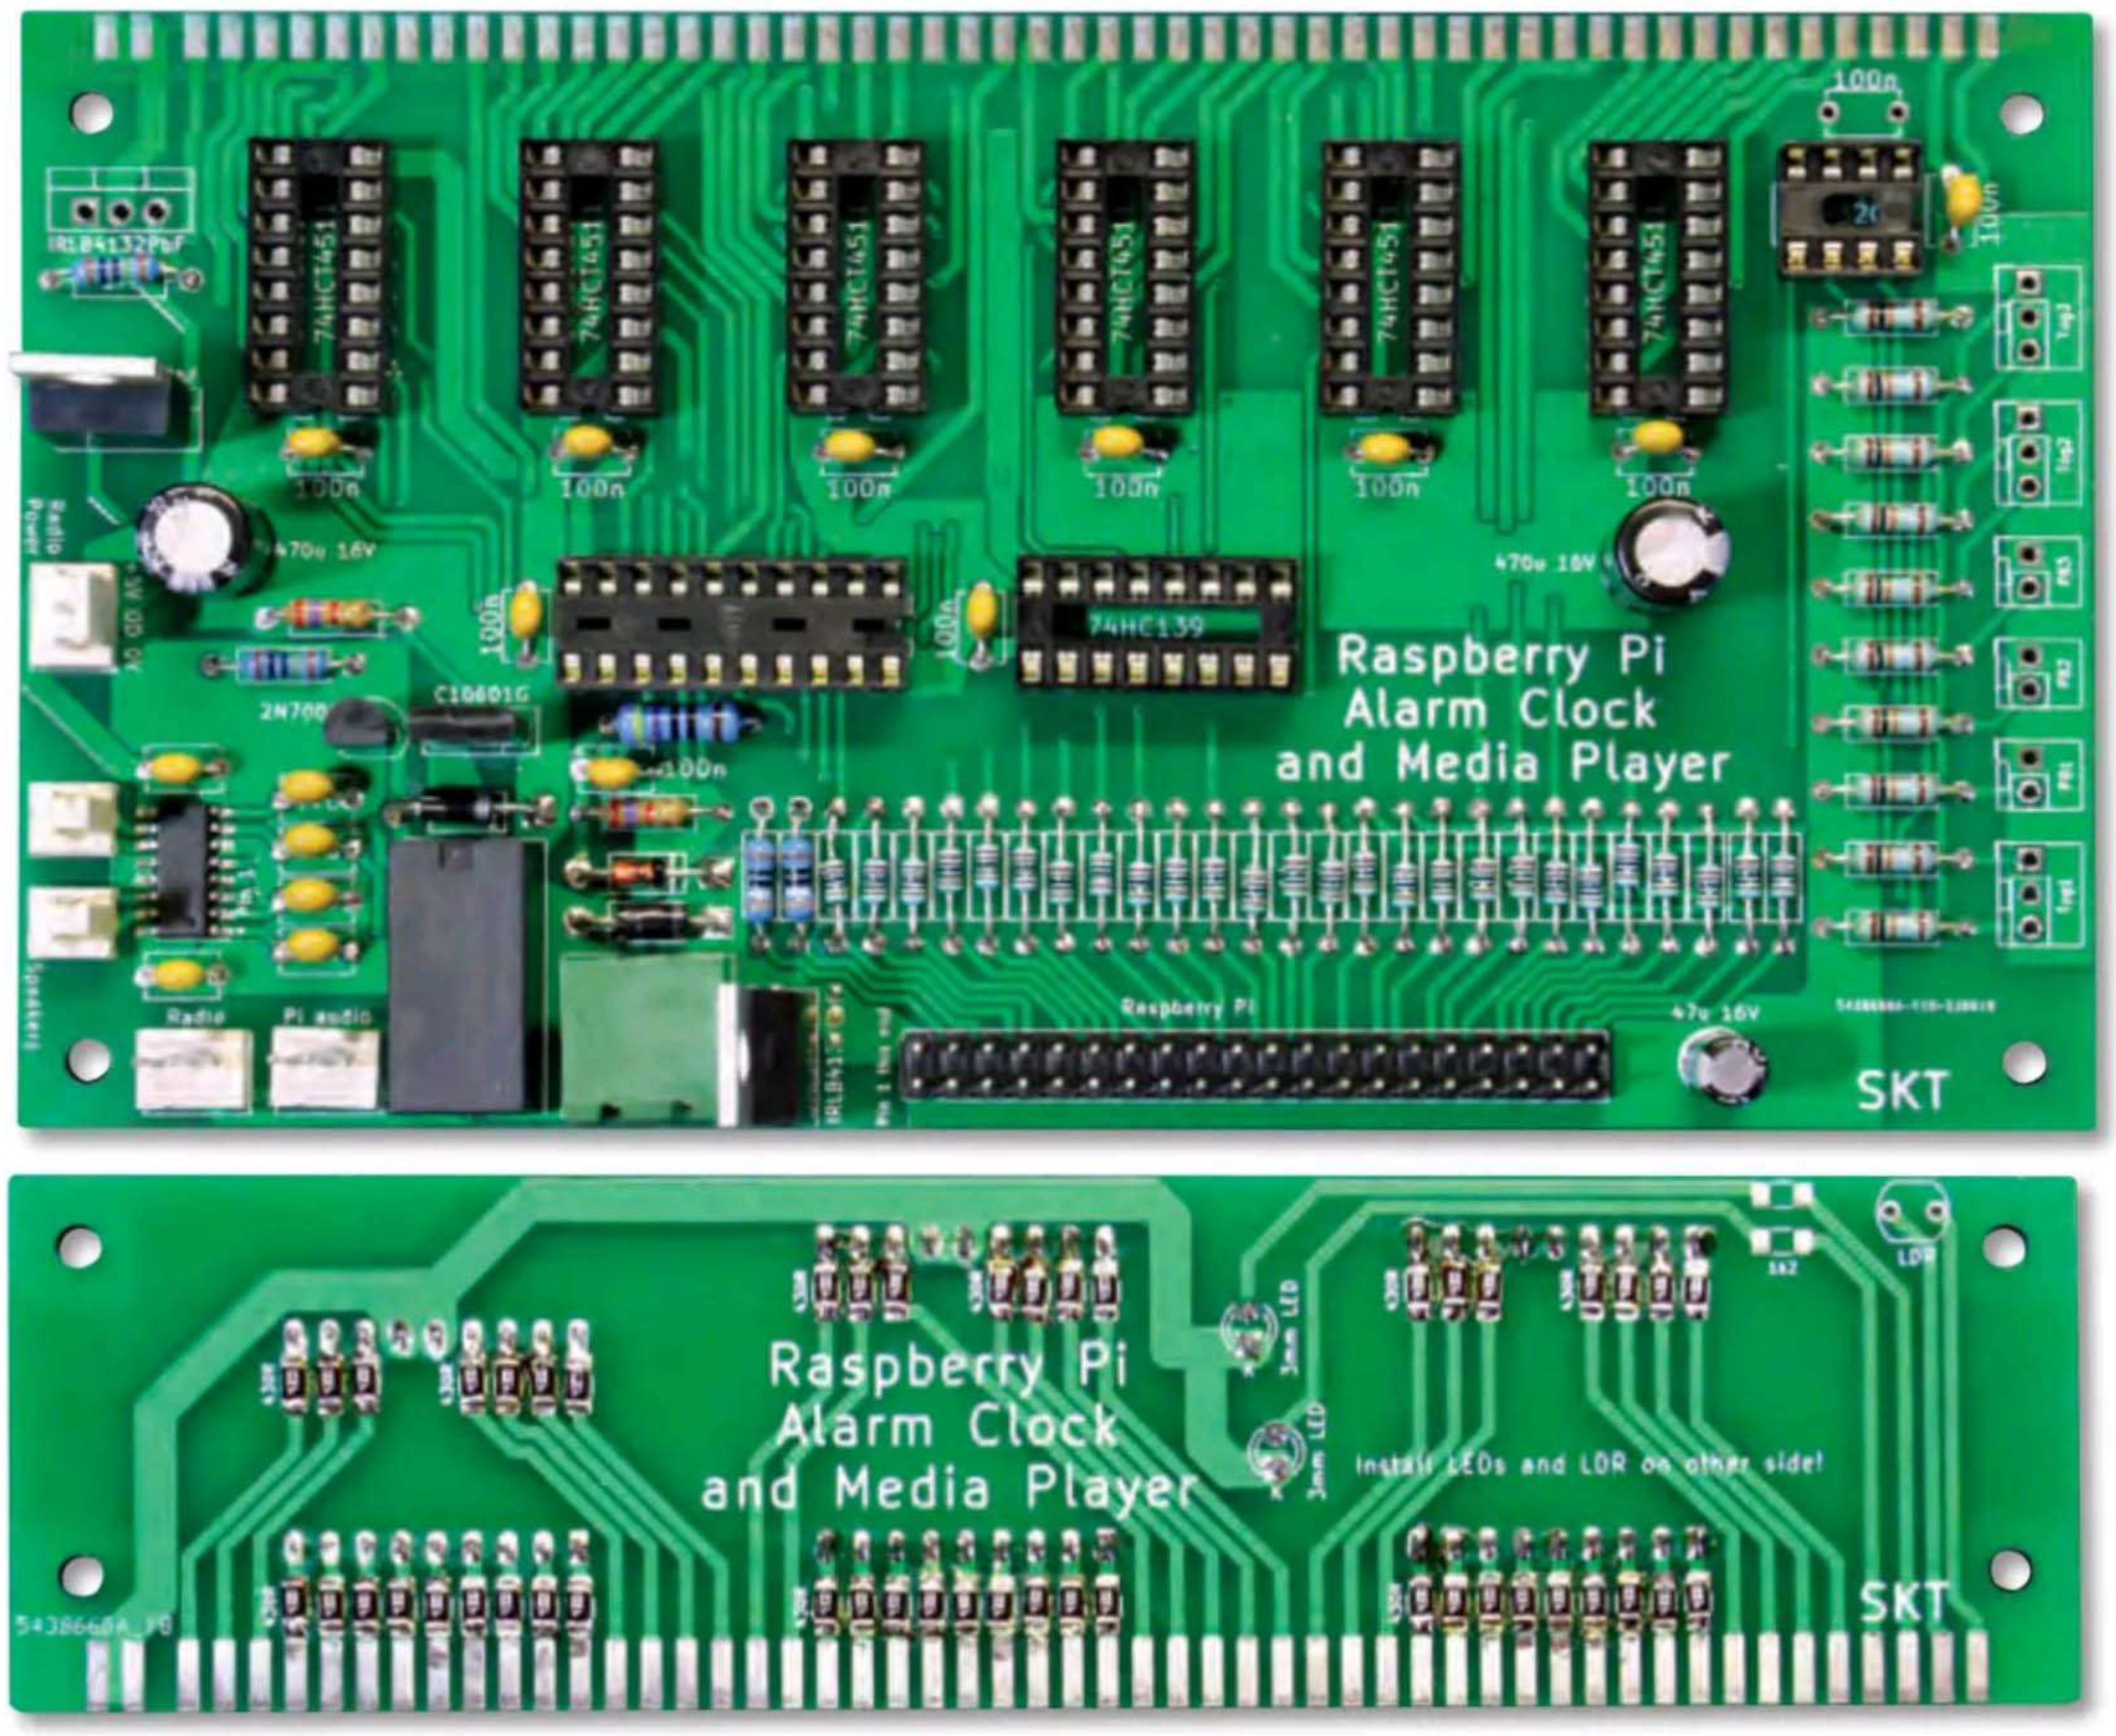
Raspberry Pi Alarm Clock and Media Player
SKT
Raspberry Pi Alarm Clock and Media Player
Install LEDs and LDR on other side!
SKT

The two PCBs for the Raspberry Pi Clock Radio are mounted perpendicular to each other and then soldered together.

Audio

The audio section includes the amplifier that drives the speakers and an audio input for an external radio. The amplifier (IC13) is a PAM8407 Class-D low-distortion filterless amplifier chip.

At typical listening volumes, it has a distortion below 0.1% across most of the audible band. It is more than adequate for a clock radio and media player and comparable with the Pi analog audio output quality.

Three GPIO pins are dedicated to putting the amplifier into and out of standby and adjusting its volume.

The audio source is selected by DPDT relay RLY1, driven by Mosfet Q4. The GPIO line that drives Q4 also operates a second Mosfet, Q3, to act as a power switch for the external audio source. The switched external power is available at three-pin header CON5. Fig.3 shows three possible ways to power an external radio from CON5.

Q3 has a maximum voltage rating of 30V so, if using an external power source, do not exceed that.

If you don't plan to integrate an external radio or audio input, you could omit Q3 and Q4, the associated resistors, PCB headers and the relay, and fit wire links to the relay pads on the PCB to connect the Pi's audio output to the amplifier permanently.

User interface

Each switch pole or button has a 10kΩ pull-up resistor to the 3.3V rail and is connected to one of the Raspberry Pi's GPIO pins that's configured as a digital input. Therefore, when a button is pressed or a switch is toggled, the corresponding pin goes low and is detected by the software.

Power supply and protection

The reverse polarity and overvoltage protection section consists of diodes D1 & ZD1, SCR1, Mosfet Q1 and associated passive components. It protects the Clock Radio from an incorrect power supply that could otherwise damage it.

D1 protects against reverse polarity by effectively short-circuiting the supply rail if power is applied with the wrong polarity.

It will get hot, but it gets the job done. A switch-mode plugpack will enter overcurrent shutdown if shorted by D1, and your Clock Radio will not power on, allowing you to discover

Parts List – Raspberry Pi-Based Clock Radio

1 instrument case, 200 × 155 × 65mm **[Velleman WCAH2505 (black)]**
1 Raspberry Pi (model 3, 4, Zero 2W or similar)
1 sheet of green acrylic/Perspex, sized and shaped for the front panel
1 double-sided PCB coded 19101241, 150 × 83mm
1 double-sided PCB coded 19101242, 150 × 44mm
1 5V DC 2A+ plugpack
1 16-33kΩ light-dependent resistor (LDR1)
    [DigiKey PDV-P8103-ND, Farnell 3168335]
3 panel-mount SPDT centre-off momentary toggle switches (S1, S5, S6)
3 panel-mount SPST momentary pushbuttons (S2-S4)
1 J104D style 5V DC coil, 2A DPDT relay (RLY1)
    [DigiKey 2449-J104D2C5VDC.20S-ND, Farnell 1652604]
1 2×20-pin header, 2.54mm pitch
1 2.5mm chassis-mounting DC barrel socket (CON1)
1 2-way right-angle pluggable terminal block, 5.08mm pitch
6 3-way, 2.54mm pitch polarised headers with matching plugs and pins
5 2-way, 2.54mm pitch polarised headers with matching plugs and pins
2 40-pin IDC line sockets
1 20-pin DIL IC sockets
7 16-pin DIL IC sockets
1 8-pin DIL IC sockets
1 panel-mount barrel socket to suit plugpack
2 red panel-mount banana socket
2 black panel-mount banana socket
1 short stereo audio cable with a 3.5mm jack plug at one end
1 15cm length of 40-way ribbon cable
1 50cm length of figure-8 speaker cable
1 1m length of 3-way ribbon cable
2 M3 × 32mm panhead machine screws
10 M3 × 6mm panhead machine screws
2 M3 hex nuts and flat washers
6 12mm-long M3-tapped Nylon spacers
2 short lengths of medium-duty hookup wire (red & black)

Semiconductors
6 74HCT4511 7-segment display driver ICs, DIP-16 (IC4-IC9)
1 74HC139 dual 2-to-4 decoder IC, DIP-16 (IC10)
1 74HCT374 8-bit parallel latch IC, DIP-20 (IC11)
1 MCP3201-CI/P 12-bit ADC, DIP-8 (IC12)
1 PAM8407DR filterless Class-D stereo amplifier IC, SOIC-16 (IC13)
3 IRLB4132PbF 30V 78A N-channel Mosfets, TO-220 (Q1-Q3)
1 2N7000 small signal N-channel Mosfet, TO-92 (Q4)
1 C106D1G sensitive-gate SCR, TO-126 (SCR1)
2 0.8in/20.3mm green dual 7-segment display, eg, LDD-C812RI or LD0805GWK
    [DigiKey 67-1473-ND, Farnell 2627654]
1 0.56in/14.2mm green dual 7-segment display, eg, LDD-C512RI or
    LD0565GWK **[DigiKey 67-1459-ND, Farnell 2627648]**
2 green diffused 3mm LEDs (LED1, LED2)
    [DigiKey 754-1609-ND, Farnell 2112096 or equivalent]
1 5.1V 1W zener diode (ZD1)
2 1N4004 400V 1A diodes (D1, D5)

Capacitors
2 470µF 16V electrolytic (2.5mm lead pitch)
1 47µF 16V electrolytic (2mm lead pitch)
2 1µF 50V (multi-layer) ceramic
4 470nF 50V (multi-layer) ceramic
10 100nF 50V (multi-layer) ceramic or MKT
1 10nF 50V (multi-layer) ceramic or MKT

● values may need to vary to match the display segment brightness.

Resistors *(all ¼W 1% axial unless noted)*
2 1MΩ
10 10kΩ
1 2.7kΩ
2 1.3kΩ SMD M3216/1206 1% ●
1 1.3kΩ
26 1kΩ
1 470Ω ½W axial
48 430Ω SMD 1206 1% ●
2 390Ω

Installing the software on a Raspberry Pi

You will need an SD card with at least 4GB capacity. Larger is fine; you can use the extra storage to hold your media library. With Raspberry Pis, the read/write speed and quality of the SD card make a difference. Cheap SD cards often perform poorly.

The SD card must be loaded with either the Debian Bullseye Lite or Debian Bookworm Lite operating systems. Debian images older than Bullseye are not suitable. The easiest way to prepare the SD card is with "Raspberry Pi Imager", freely available for Windows, macOS and Linux.

Launch Raspberry Pi Imager, insert the SD card into your computer (via a card writer if it doesn't have a slot) and click the CHOOSE DEVICE button, then select "No Filtering". For a Pi 4 or Pi Zero 2W, choose Raspberry Pi OS (Other) → Raspberry Pi OS Lite (64-bit). For other models, select Pi Raspberry OS (Other) → Raspberry Pi OS Lite (32-bit). Then click CHOOSE STORAGE to select the SD card, click NEXT, pick EDIT SETTINGS and fill out the form:

1. Set a unique hostname for your clock ("clock" if you can't think of anything else).
2. Enable SSH using password authentication.
3. Set a username and password for logging in via SSH.
4. Enter your wireless LAN details (SSID, password and country).
5. Set the locale settings for your area.
6. Deselect the option to eject media (the SD card) when finished, as you'll also be copying the clock software to the SD card before 'ejecting'.

Write down the hostname, username and password so you can log into the Pi later. Next, click SAVE, then YES then WRITE. When the card has been written, download the clock software zip file from the Silicon Chip website.

Inside the zip file is a file named "alarm-clock_v01.tgz" that you need to copy onto the SD card. Copy the TGZ file from the ZIP archive to the root of the "bootfs" directory on the SD card the same way you transfer files to a thumb drive. The v01 number could increase in future if there are updates to the software. Finally, eject the SD card, insert it into the Pi and apply power. The ZIP archive also contains a PDF document with screenshot of the installation, and post-publication notes.

Connecting to the Pi

Because there's no video output, the only way to know the Pi is ready to proceed is to connect to it over your network (wired or WiFi). The first time a Pi boots, it could take a few minutes longer than usual. To avoid frustration, apply power and make a cup of tea or coffee.

You will need an SSH client to connect to the Pi. In Windows, you can use PuTTY or OpenSSH; macOS and Linux have 'ssh' command line tools. You can connect using its IP address or the hostname specified when you prepared the SD image.

Most home routers generally publish local hostnames using a ".local" suffix, as suggested in Raspberry Pi Imager. So you can try to connect to "clock.local" (or whatever other name you chose). If that does not work, consult your router's documentation or look at the router's DHCP leases table to find the IP address allocated to the Pi.

When you connect, the Pi will prompt for the username and password that you specified during the SD card setup. Enter them to log in and get the remote command prompt.

Finishing the clock software installation

On the Pi, the file you copied to the SD card earlier is available within the bootfs partition at /boot. You can now extract the contents using the command:

```
tar zxf /boot/firmware/alarm-clock_v01.tgz     ← for Bookworm OS
tar zxf /boot/alarm/alarm-clock_v01.tgz        ← for Bullseye OS
```

This command creates a subdirectory called "alarm-clock" containing the source code and will also leave an installation script in your current directory. The last stage in the software installation is to run that installation script (you must copy this exactly, including the letter case):

```
sudo ./Install_Clock.sh
```

The installation script asks for your password twice, whether you would like to install firewall rules that prevent access from IP addresses originating on a different subnet (you will probably want to say yes) and then asks if you would like to attempt to disable the power and activity LEDs.

Web-based configuration

To reach the web interface, open a browser and surf to http://clock.local or whatever system name or IP address you used to ssh into the clock. You're greeted by the clock's home page, which contains links to the various configuration and media player functions, a summary of the configured alarms, the playlist if media is currently playing, and a list of any other clocks found on the local network.

We'll have more information on configuring the clock in part two next month, along with instructions on updating the software, using it as a Bluetooth speaker, testing and more information on the clock software.

If you run into trouble during installation you should check the instructions included with the software download, as steps may have changed after publication. These instructions are for the 1.8.1 version of Raspberry Pi Imager, but earlier versions will work with slight changes.

the mistake without losing any smoke.

The over-voltage protection isolates the rest of the circuit from the supply if the supply voltage exceeds about 5.7V. With a normal supply of around 5V, zener diode ZD1 does not conduct, so the gate of SCR1 remains at 0V. The 2.7kΩ pull-up resistor pulls the gate of Mosfet Q1 up to +5V, switching it on and connecting circuit ground to the incoming supply's negative terminal.

If the supply voltage exceeds 5.7V, there is around 0.6V at the gate of SCR1, so it switches on, pulling the gate of Mosfet Q1 to 0V. That switches Q1 off, allowing the circuit ground to rise to the positive supply rail, leaving no voltage to power the remainder of the circuit. The potential for damage to the Pi starts at around 6V, so the SCR trigger voltage is just slightly below that.

SCRs behave a little like bipolar NPN transistors acting as switches, except that SCRs latch themselves on after their trigger voltage has been reached. This way, Q1 remains off until the offending power supply is disconnected, at which point it resets.

Component selection

When purchasing components for this project, note that electrolytic capacitors come in all shapes and sizes. The hole spacing for the two 470µF electrolytics is 2.5mm, while the 47µF electrolytic holes are spaced at 2mm. Most 16V rated capacitors will have similar lead spacings but higher-voltage electros may not fit well.

If possible, measure the actual component or check the catalog or data sheet to find a good match. The clock will work with higher-voltage or larger components, but they may not fit as neatly on the board.

Sockets are recommended for the DIP ICs. If ever you need to replace a chip, extracting the IC from a socket will be much easier than desoldering it from the joined main and display board assembly.

However, sockets can slowly oxidise over time and eventually cause problems; soldered chips are generally more reliable in the long term. Removing the chip from its socket and then reinserting may be all that's required to re-establish good contact.

The second article next month will have all the construction details, usage instructions and information on updating the firmware. PE

Max's Cool Beans

By Max the Magnificent

Arduino Bootcamp – Part 24 (it's time to get serial)

You know how people sometimes tell you, "I've got good news and bad news, which do you want first?" Don't worry. I'm not going to say that. What I will say is, "I have happy news and sad news".

Also, I'm not going to give you a choice because we'll start with the opposite of happy news, which is that this is to be the last in our current collection of *Arduino Bootcamp* columns (😟).

Meanwhile, the opposite of sad news is that we are going to start a new series of *Weird and Wonderful Arduino Projects*, which will use many of the techniques we've explored together, including some we are about to discover in this column (🙂). But first…

B or NotB, that is the question

Before we dive headfirst into the fray with gusto and abandon, following our previous column in the November issue, several readers emailed me, moaning and groaning that the puzzle I posed is unsolvable. "Oh, ye of little faith" is all I can say. Well, not quite; I follow this with, "Watch and learn, Grasshopper".

As you may recall, and as illustrated in Fig.1(a), our problem involves a black box with three inputs (`A`, `B`, `C`) and three outputs (`NotA`, `NotB`, `NotC`). Each output is to be the logical negation of its corresponding input, so a 0 on `A` results in a 1 on `NotA`, and vice versa; the situation for the other input-output pairs is similar.

We are permitted by those who don the undergarments of authority to use only two NOT gates along with as many 2- and 3-input AND and OR gates as we wish.

There are various ways in which we can tackle this problem. The solution shown in Fig.1(b) was generated by my friend George Harper. Personally, I think George's resolution to this conundrum is awe-inspiring in its elegance.

I should note that George is using a pseudo-language of his own devising to represent this solution (most computer languages won't allow us to start a variable name with a numerical character, for example). Also, George employs '&', '|', and '!' characters in his Boolean equations to represent AND, OR, and NOT functions, respectively.

This really is cool. Once we've specified the names of our input and output signals, we move onto the Boolean equations describing the internal nodes. These equations can be transferred one-for-one into corresponding logic gates.

In the following discussions, "#n" refers to the line number comments associated with the internal nodes. All of this is pretty much self-explanatory. Remember that 0 and 1 equate to the Boolean concepts of `false` and `true`, respectively.

This means we can read equation #1 as saying, "internal node `2or3ones` will be 1 (`true`) if any two or all three of inputs `ABC` are 1; otherwise, it will be 0 (`false`)." The physical implementation of this expression requires three 2-input AND gates and one 3-input OR gate.

Similarly, equation #2 tells us, "the value of internal node `0or1ones` is NOT the value of internal node `2or3ones`". Another way saying that is, "internal node `0or1ones` will be 0 if `2or3ones` is 1; otherwise, it will be 1".

An alternative approach to wrap our brains around this part of the problem is that there are only four broad possibilities: all the inputs are 0, one of the inputs is 1, two of the inputs are 1, or all three of the inputs are 1.

Since we know that `2or3ones` will be 1 if two or three of the inputs are 1, it will be 0 only if none or only one of the inputs are 1.

Since we generate our `0or1ones` value simply by inverting whatever value is present on `2or3ones`, the implementation of this expression will only require one of our two NOT gates ("our preciouses").

As we can see, this solution is layered. Equation #3 means, "the internal node `1one` will be 1 if both `0or1ones` is 1 and **any** of the `ABC` inputs are 1". The physical implementation of this expression will require only a 2-input AND gate and a 3-input OR gate.

Similarly, Equation #4 tells us, "the internal node `1or3ones` will be 1 if `1ones` is 1 or **all** of the `ABC` inputs are 1". The physical implementation of this expression requires only a 2-input OR gate and a 3-input AND gate.

Equation #5 is similar in concept to equation #2. Since the only possibilities are for 0, 1, 2 or all 3 of the inputs to be 1, then if `1or3ones` is 1 (`true`), `0or2ones` must be 0 (`false`), and vice versa. We can achieve this by inverting the output of `1or3ones` with our second and final NOT gate.

Using equations #6 and #7 to generate the remaining internal nodes, `0ones` and `2ones`, is easy peasy lemon squeezy. The physical implementation of each requires just a single 2-input AND gate.

The key takeaway from all this is that we now have three extremely useful signals: `0ones`, `1one`, and `2ones`, which tell us if none, only one, or

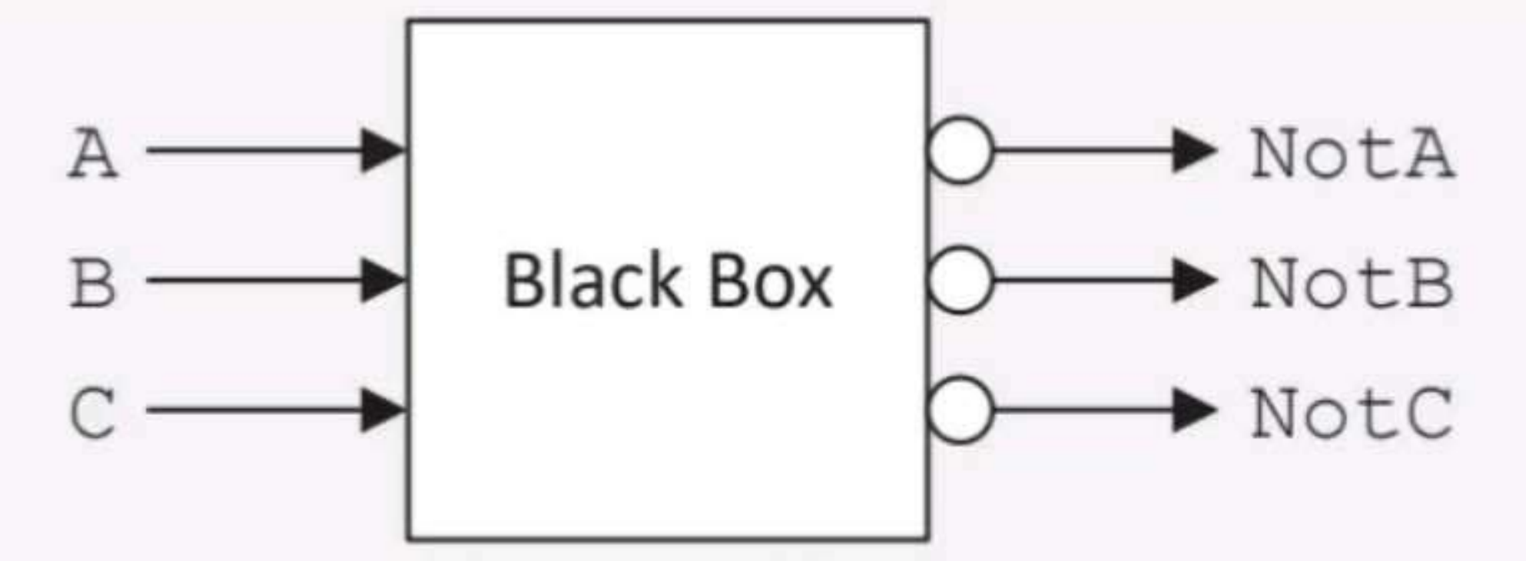

(a) The problem

Inputs:

```
A, B, C;
```

Outputs:

```
NotA, NotB, NotC;
```

Internal Nodes:

```
2or3ones = ( (A & B) | (A & C) | (B & C) ); // #1
0or1ones = !(2or3ones);                     // #2
1one     = 0or1ones & (A | B | C);          // #3
1or3ones = 1one | (A & B & C);              // #4
0or2ones = !(1or3ones);                     // #5
0ones    = 0or2ones & 0or1ones;             // #6
2ones    = 0or2ones & 2or3ones;             // #7
```

Equations for Outputs:

```
NotA = 0ones | (1one & (B | C)) | (2ones & (B & C));
NotB = 0ones | (1one & (A | C)) | (2ones & (A & C));
NotC = 0ones | (1one & (A | B)) | (2ones & (A & B));
```

(b) The solution

Fig.1: proof that there is a solution to the problem!

only two of our three ABC inputs are 1, respectively. Having access to these signals means we now have sufficient information at our disposal to generate our three output values.

This is where things get so clever I'm tempted to squeak with excitement. We want an output to be 1 only if its corresponding input is 0. Pick any output.

This output will be 1 if `0ones` is 1 (meaning all the inputs are 0), or if `1one` is 1 and **either** of the other inputs is 1 (meaning this input is 0), or if `2ones` is 1 and **both** the other inputs are 1 (again meaning this input is 0). This is reflected in our three output equations.

Every time I look at this solution, a smile comes to my face, and I think, "Wow! That. Is. So. Ingenious!"

Pondering our persisting poser

Do you remember the other puzzle I posed in last month's column? I said to assume our four 7-segment light-emitting diode (LED) displays are only ever going to be used to implement a clock. Also, to assume that this clock can support both 12- and 24-hour modes and that we are suppressing any leading zeros in the most-significant hours digit (we introduced this concept in *PE*, August 2024).

Although we didn't explicitly state it, let's also assume that we are going to exclude the displays' decimal point segments from these discussions.

Based on all these assumptions, what would be the maximum number of segments we might see active at the same time? And what would be the ideal value for our current-limiting resistors to obtain the maximum possible segment current values, assuming a total allocation of 350mA to power all four of our displays?

I bet this is not as easy as you think. If you haven't already done so, why don't you take a stab at this now, and then compare it to my answer, which I'll present at the end of this column.

Lest we forget

As per our previous column, the current implementation of our clock employs four latching CD4511 BCD (binary-coded decimal) to 7-segment decoders. This means each of our 7-segment displays has its own BCD decoder, each of which latches (stores) the BCD value we wish it to display.

As usual, you can download a PDF of our existing ensemble (named **CB-Dec24-brd-01.pdf**). And, as usual, all the files mentioned in this column are available from the December 2024 download page on the *PE* website: https://pemag.au/link/ac1m

Since we are no longer multiplexing the displays, they maintain a constant brightness, irrespective of the number we decide to use.

However, there are two disadvantages to this scenario. The first is that, in addition to the 4 `ABCD` signals driving the inputs to the BCD decoders, each decoder also requires its own `LE` (latch enable) signal.

Since we are currently using four displays, this means we are consuming 4 + 4 = 8 of our Arduino Uno's 12 precious digital input/output (I/O) pins. Every time we add a digit, we will lose another pin (😔).

The second disadvantage is that BCD decoders don't control the `dp` (decimal point) segments on 7-segment displays. As we will come to see, these segments can be employed to convey all sorts of useful information.

So, if we decide we want to use our `dp` segments, we will have to control them individually, which means losing another pin for each display. In the case of our current 4-digit clock, this means we will end up consuming 4 + 4 + 4 = 12 digital pins, leaving us with no pins available to implement other functions, like driving a buzzer, for example.

Shift over

Happily, there is another solution open to us. One that will potentially allow us to drive as many 7-segment displays as we wish—including their `dp` segments—using only 3, 4, or 5 of the Arduino's digital I/Os, depending on our implementation.

In an earlier column (*PE*, September 2024), we introduced the concept of memory elements (or 'registers') called D-type flip-flops. We also described how these little scamps could be connected to form a serial in, parallel out (SIPO) shift register.

Suppose we form an 8-bit 'bare bones' SIPO shift register as illustrated in Fig.2(a). All the registers share a common clock signal, which we've called `SRCLK` (shift register clock). The data input to the first register in the chain is called `SER` (serial data in). The output from each register drives the next register in the chain.

When a rising edge (a 0 to 1 transition) is applied to `SRCLK`, the value on `SER` will be loaded into the first register. At the same time, the existing contents of the first register will be loaded into the second, the contents of the second register will be loaded into the third, and so on down the chain.

One way to visualize this is as a line of eight people passing objects to each other when someone shouts "Now!".

So, to load an 8-bit value into our shift register, we first place the first a 0 or 1 on the `SER` input, then apply a positive-going pulse to the `SRCLK` input. We then place the next 0 or 1 on the `SER` input, apply another positive-going pulse to the `SRCLK` input, and so forth until all eight bits have been loaded.

This means we need to load the last bit first and the first bit last, if you see what I mean.

The outputs from all eight registers, named Q_A through Q_H in our illustration, are presented to the outside world, where they can be used to drive all sorts of things like… well, like a 7-segment display.

In our case, we are going to use the Q_A through Q_G outputs to drive the corresponding `a` through `g` segments on our display, while the Q_H output can be used to control our `dp` segment.

Now, this is the clever bit. Assume our 8-bit shift register is presented and packaged as an integrated circuit (IC). We can daisy chain multiple devices together to form a longer shift register. In this case, the `SRCLK` signal would be used to drive all the chips. Meanwhile, the Q_H output from the first chip could be used to drive the `SER` input to the second device, and so on down the line.

Returning to Fig.2(a), observe that we've taken a copy of the Q_H output that we've called $Q_{H'}$. It's this $Q_{H'}$ output that we will use to drive the next device in the chain. The reason for our doing this will become apparent shortly.

Just to keep things simple, we will refer to these outputs more concisely as `abcdefgh` in the discussions below.

The main point to note at this stage is that, by daisy-chaining a bunch of these devices together, we can now control as many displays as we wish using only two of our Arduino's digital I/O pins. Pretty cool, eh?

All clear!

Suppose we decide to daisy chain four 8-bit SIPO shift registers and use them to drive our four 7-segment displays. Further suppose we wish to clear the registers, loading each element with a logic 0 value. In the case of the bare-bones implementation shown in Fig.2(a), we could achieve this by presenting a 0 to the `SER` input of the first register, then applying 32 clock pulses to the `SRCLK` input.

Our Arduino Uno's system clock is running at 16MHz (16 million cycles a second), so loading our shift register with 0s will take only an imperceivable fraction of a second.

However, since we often wish to clear the contents of our shift registers, and some applications will involve very long shift register chains, it's common to provide a $\overline{\text{SRCLR}}$ (shift register clear) input, as illustrated in Fig.2(b). It's also

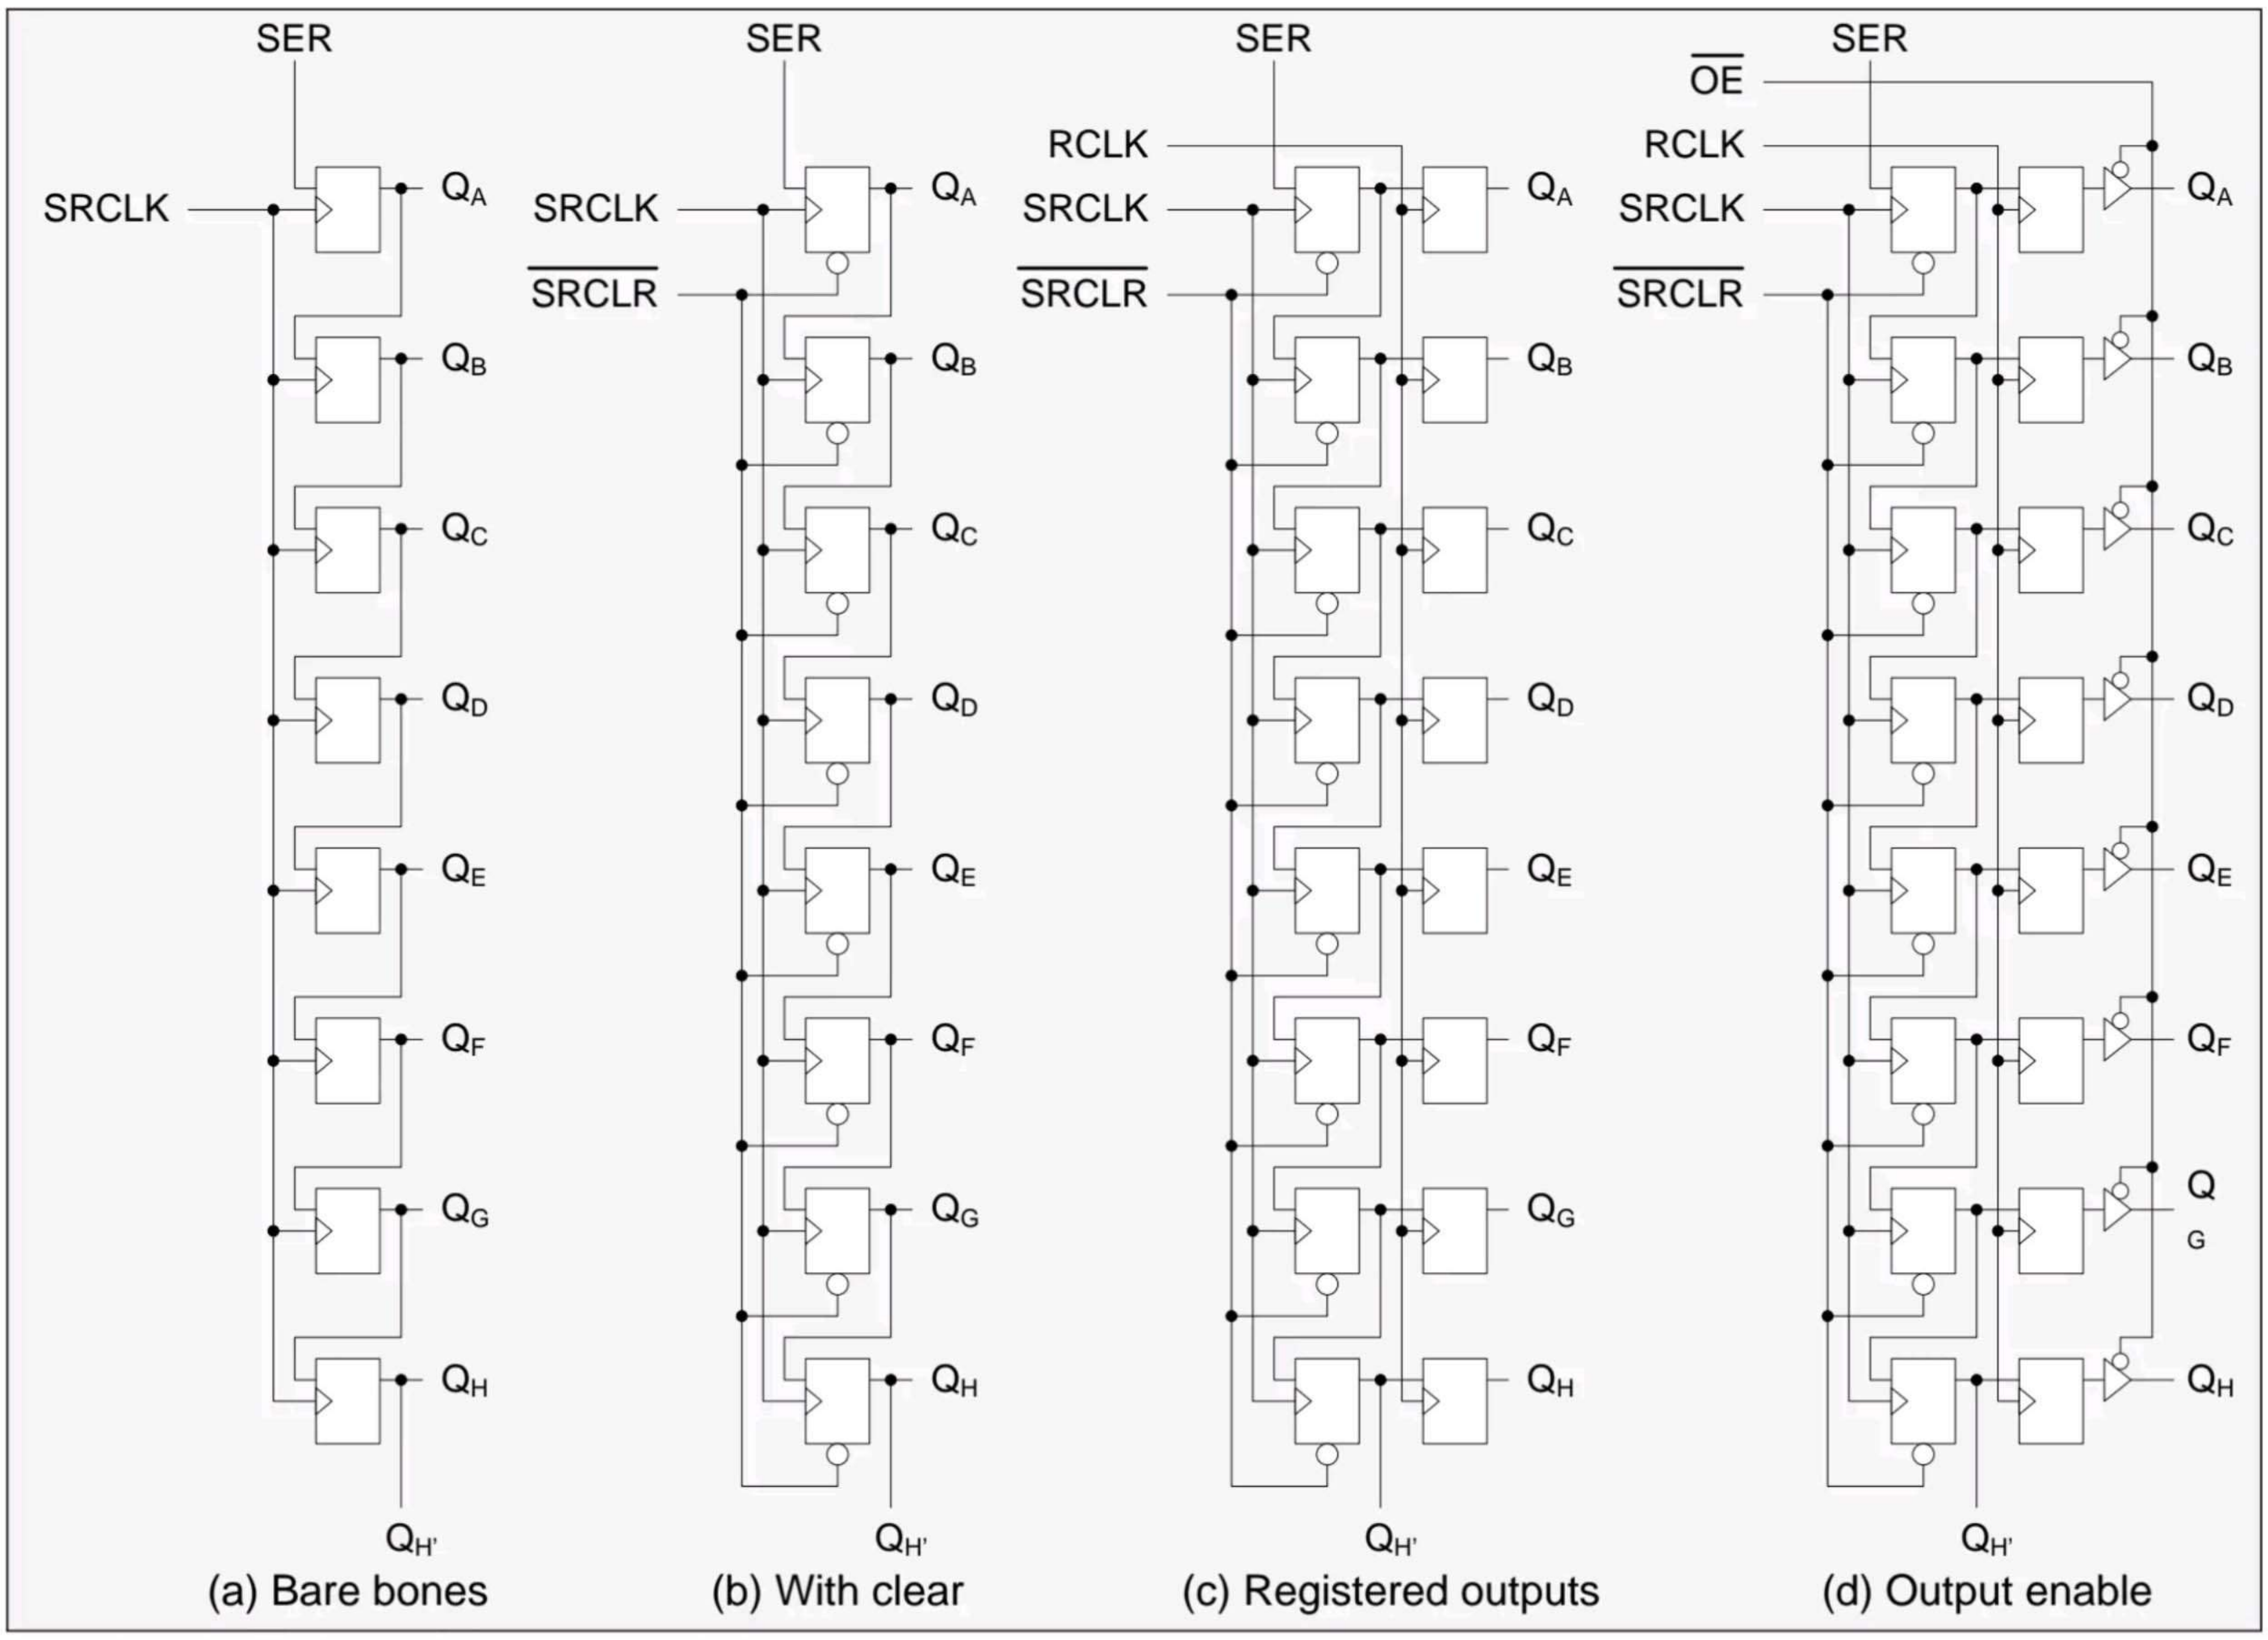

(a) Bare bones (b) With clear (c) Registered outputs (d) Output enable

Fig.2: four different flavours of eight-bit SIPO (serial in, parallel out) shift register.

common for this signal to be active-low, as shown by the bar over its name.

Once again, the $\overline{\mathtt{SRCLR}}$ signal from our microcontroller will be connected to all the shift registers in the chain.

A fly in the soup

Unfortunately, there's a fly in the soup and an elephant in the room (I never metaphor I didn't like). Suppose we connect the outputs from our 8-bit shift register to a 7-segment display (via current-limiting resistors, of course).

Now suppose we wish to display the number '4'. This requires segments `ade` to be turned off, while segments `bcfg` will need to be turned on as illustrated in Fig.3.

Remember that we are planning on using the `abcdefg` and `h` outputs from our shift register to drive the `abcdefg` and `dp` segments on our display, respectively. So, to display '4', we need a value of 01100110 in our shift register. The problem is that we need to shift this value in through the `SER` input using eight clocks (Fig.4).

Just for giggles, let's assume we power up with our shift register containing random/unknown 0 and 1 values (shown as red question marks in Fig.4). While we are in the process of clocking our new values in, our display will respond as shown in Fig.5.

All this will take place in a fraction of a second. If you blink, you'll miss it, which means this may not be particularly important if we wish to update our display relatively infrequently.

However, if we are constantly loading new values into our shift register, even if the values are the same (like repeatedly displaying '4', '4', '4'…), we could end up with blurred, 'ghosting' effects on our display (the concept of ghosting was introduced in *PE*, July 2024).

The solution is to add an additional set of D-type flip-flops as illustrated in Fig.2(c). These output registers are controlled by their own clock, which we've called `RCLK` (register clock). We are showing `RCLK` as being positive-edge triggered, so these registers will be loaded when a rising edge (0 to 1 transition) is applied to `RCLK`.

The idea here is that the output registers will maintain whatever 0 and 1 values we last loaded into them, and it's these values that are presented to the outside world. We can load a new series of 0 and 1 values into our shift register using 8 pulses on `SRCLK` without affecting the contents of the output registers.

Once we are ready to rock and roll, a single pulse on `RCLK` will copy these new values into the output registers As usual, the `RCLK` signal from the

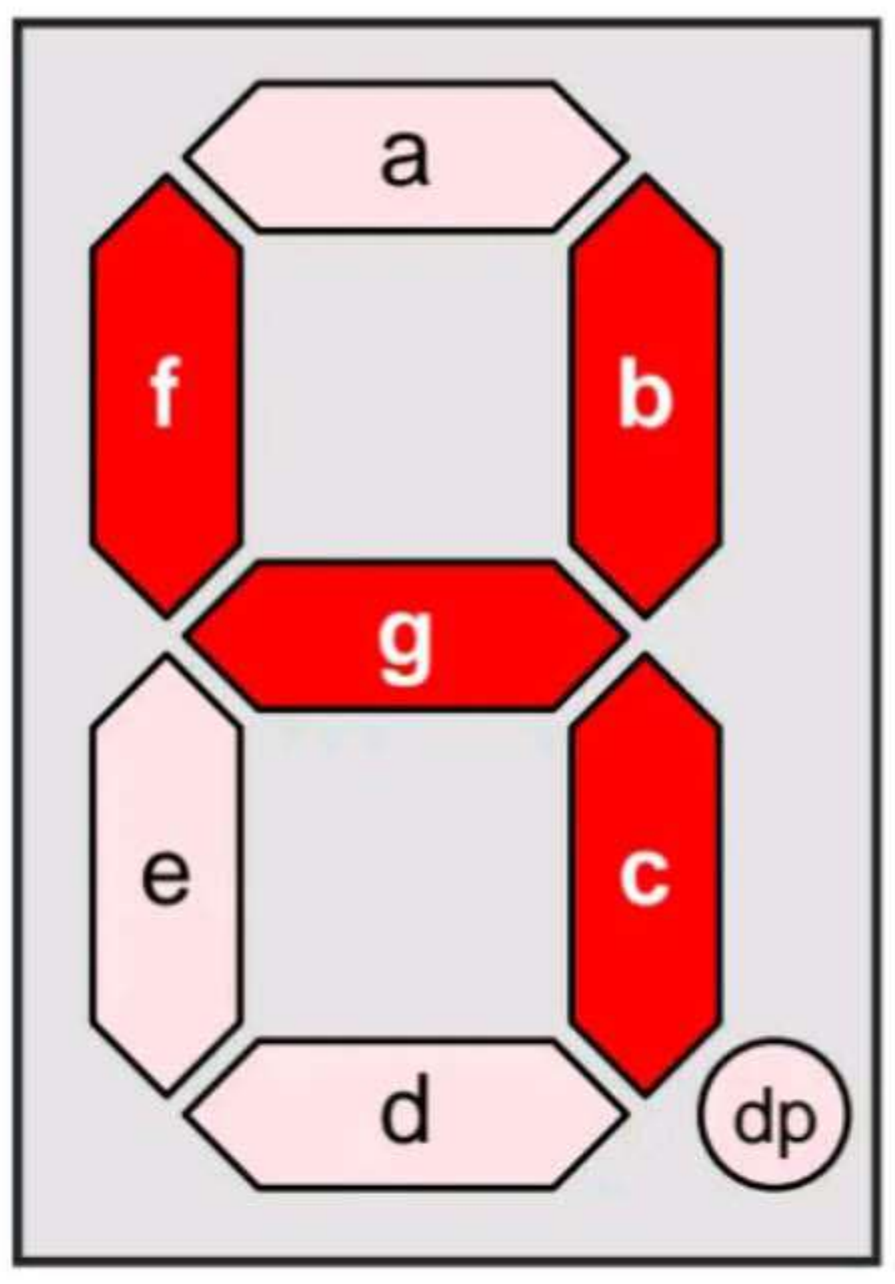

Fig.3: displaying a '4'.

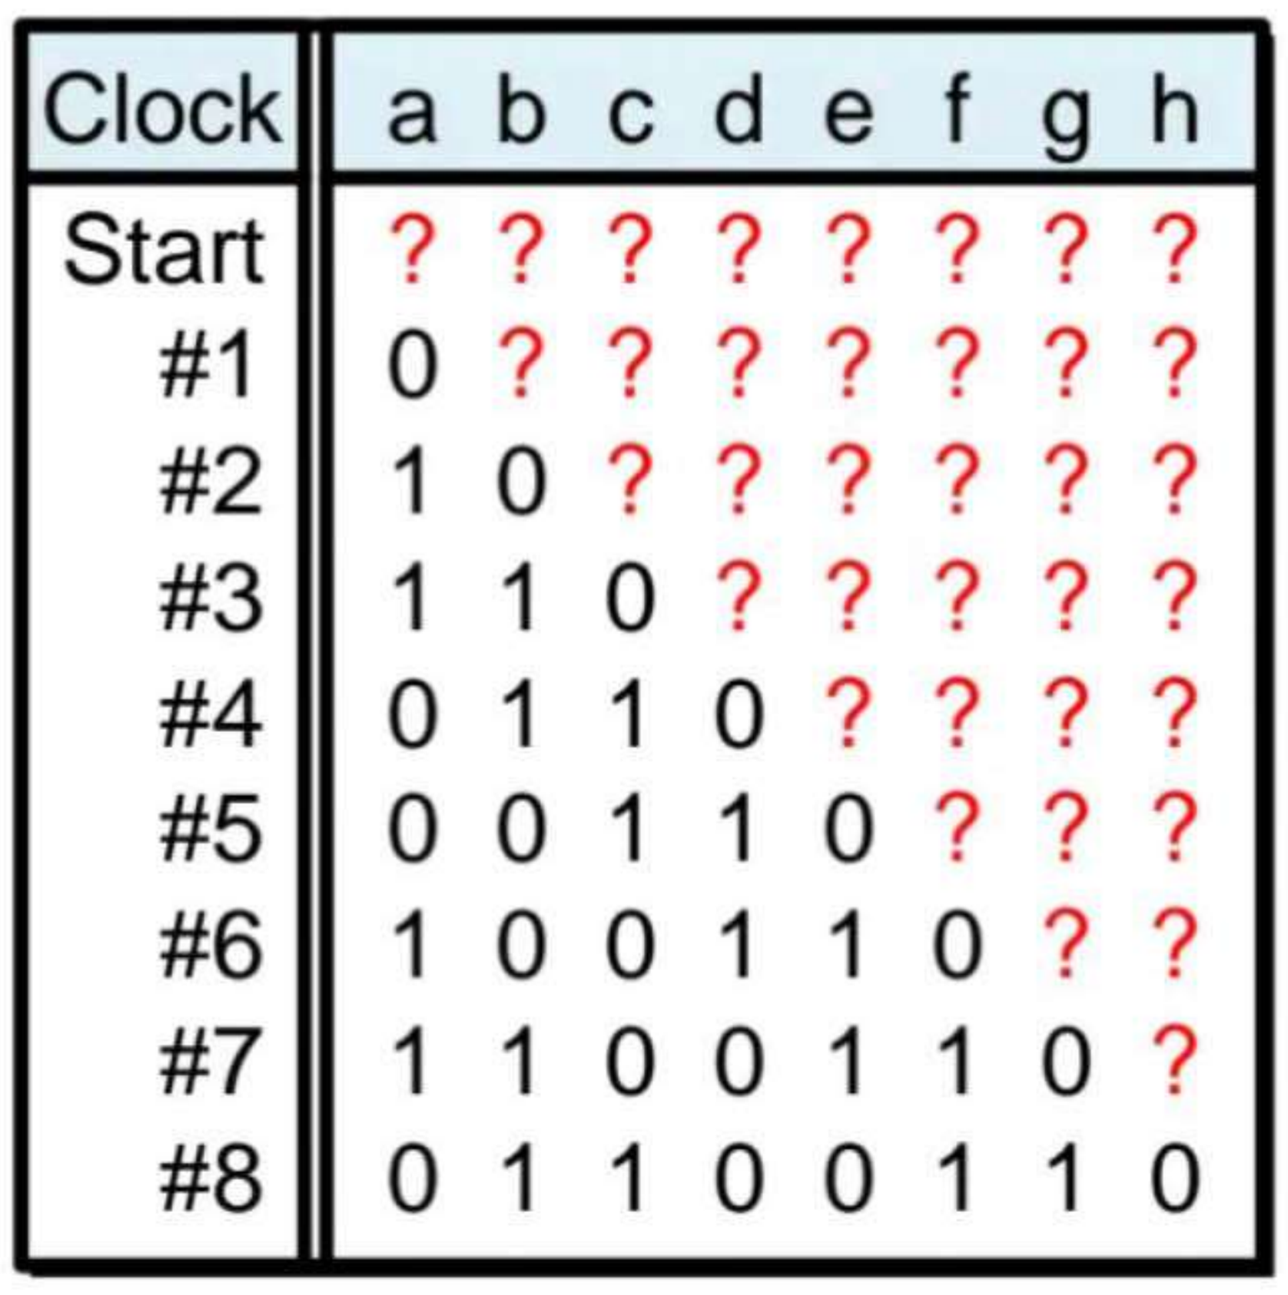

Clock	a	b	c	d	e	f	g	h
Start	?	?	?	?	?	?	?	?
#1	0	?	?	?	?	?	?	?
#2	1	0	?	?	?	?	?	?
#3	1	1	0	?	?	?	?	?
#4	0	1	1	0	?	?	?	?
#5	0	0	1	1	0	?	?	?
#6	1	0	0	1	1	0	?	?
#7	1	1	0	0	1	1	0	?
#8	0	1	1	0	0	1	1	0

Fig.4: loading our 8-bit value.

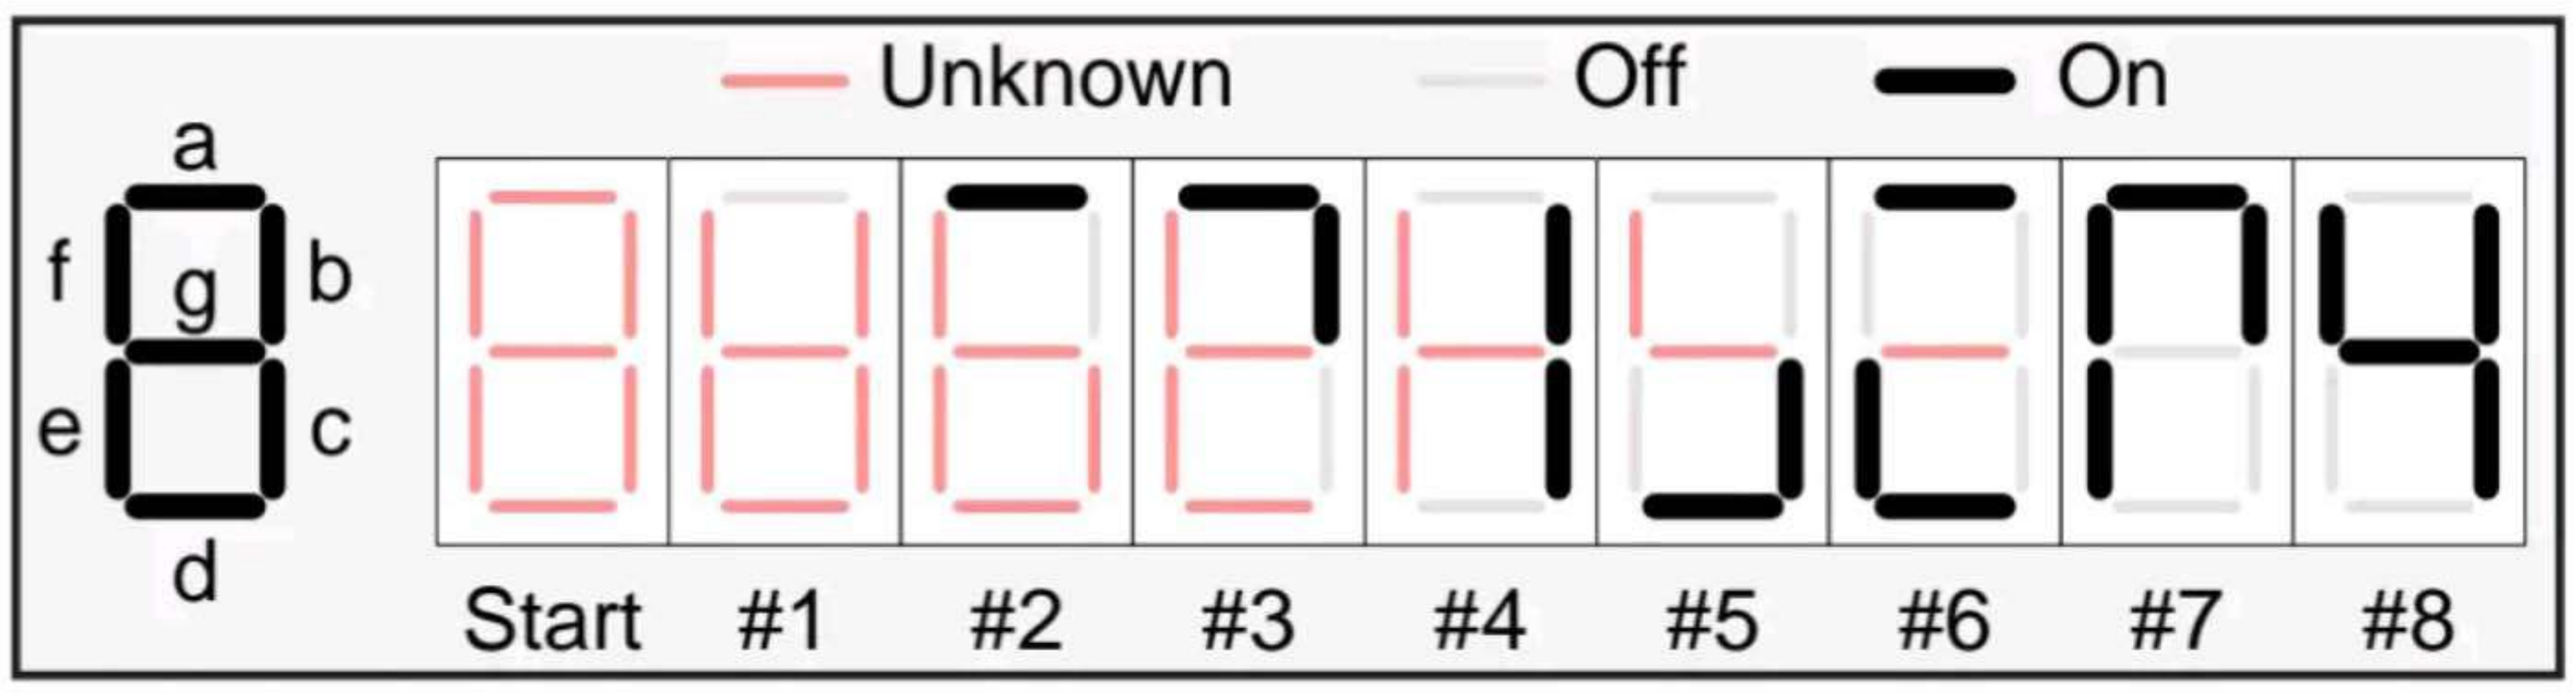

Fig.5: how the display changes while loading the new data.

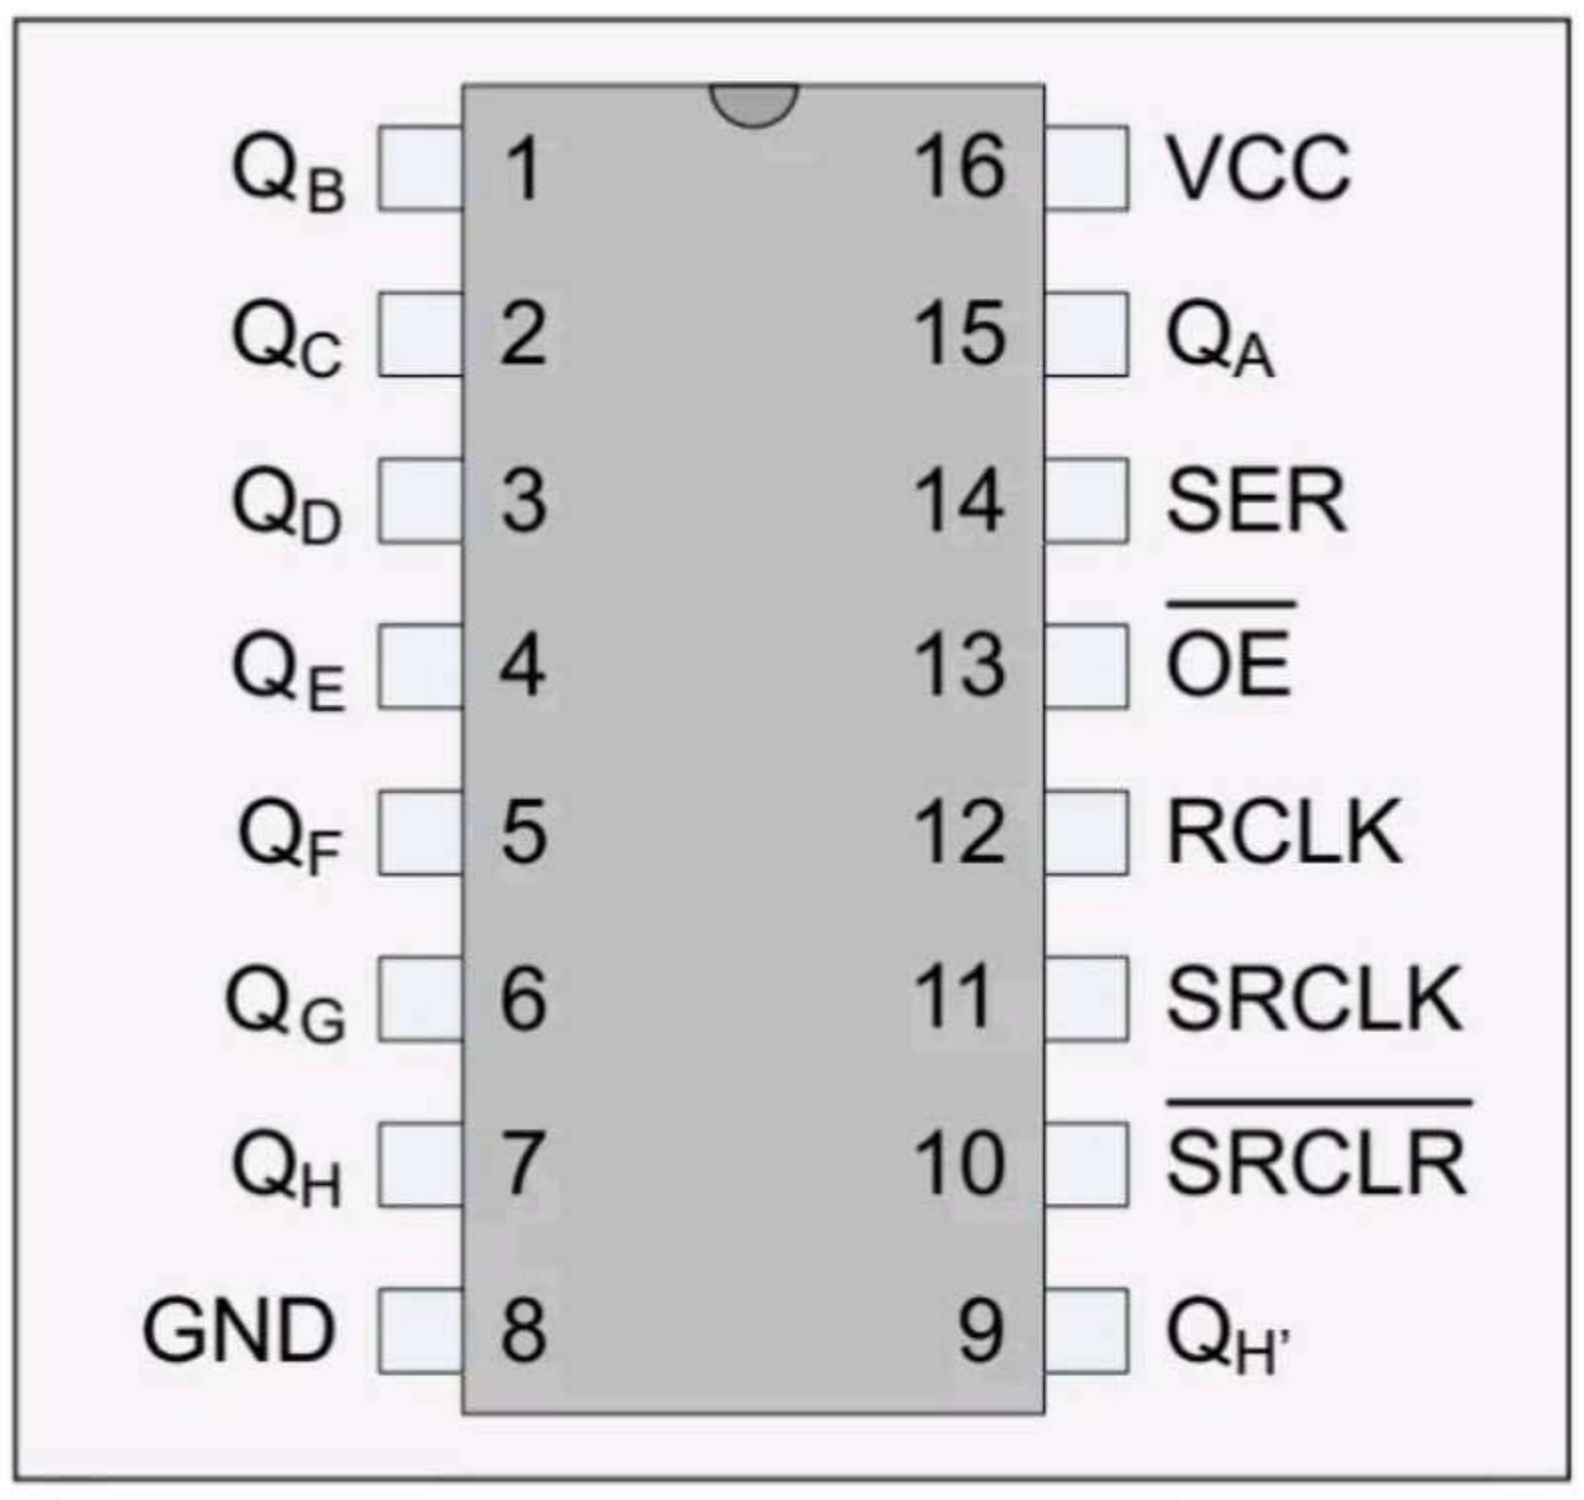
Q_B 1
Q_C 2
Q_D 3
Q_E 4
Q_F 5
Q_G 6
Q_H 7
GND 8
16 VCC
15 Q_A
14 SER
13 $\overline{\text{OE}}$
12 RCLK
11 SRCLK
10 $\overline{\text{SRCLR}}$
9 $Q_{H'}$

Fig.6: top view of the 74xx595 shift register.

microcontroller will be connected to all the shift registers in the chain.

Chaining all the devices using their $Q_{H'}$ outputs means we can update all their contents and copy their new values into the output registers simultaneously.

Feeling enabled

One more feature that is commonly found in devices like shift registers is the ability to disable their outputs. Consider the addition of the active-low $\overline{\text{OE}}$ (output enable) input as illustrated in Fig.2(d).

When this signal is in its active (0) state, whatever values are stored in the output registers will appear on the device's `abcdefgh` outputs. In contrast, when this signal is in its inactive (1) state, the `abcdefgh` outputs will be placed in a tristate condition. That effectively disconnects what they are connected to (our 7-segment displays).

The $\overline{\text{OE}}$ signal from the microcontroller will be connected to all the shift registers in the chain. In our clock project, we could employ this signal to control the brightness of the displays by using pulse-width modulation (PWM) to turn all the outputs from the shift registers on and off very quickly. We introduced the concept of PWM in *PE*, November 2023).

There's a chip for that

I bet you're wondering if there's a device that fulfills all the requirements exhibited by Fig.2(d). Well, wonder no more, because the 74HC595 ticks all the checkboxes. I had a quick Google while no one was looking and found a five-pack on Amazon for £6, plus £9 shipping (**https://pemag.au/link/ac1n**), but there are lots of other sources.

When we look at the data sheet for this little rascal (**https://pemag.au/link/ac1o**), we discover that they've used the same pin names as us (Fig.6). What are the odds? It's spooky, isn't it!

We aren't obliged to use all the features offered by this device. A minimal implementation would involve us using only three microcontroller pins to drive the `SER`, `SRCLK`, and `RCLK` pins on the shift register(s).

In this case, we would connect a pull-up resistor to the $\overline{\text{SRCLR}}$ pin to place it in its **inactive** state (we can clear the register by loading it with a string of 0s), and a pull-down resistor to the $\overline{\text{OE}}$ pin to hold it in its **active** state.

Alternatively, we might opt for a maximal implementation in which we also use the $\overline{\text{SRCLR}}$ and $\overline{\text{OE}}$ pins, thereby requiring us to dedicate five pins from our microcontroller.

For the purposes of this column, I chose to straddle the fence by using four pins to control my shift register chain (I decided to reserve the $\overline{\text{OE}}$ feature for another day).

Clearing the decks for action

Our new circuit diagram using four 74xx595 shift registers is shown in Fig.7. Pin 2 from the Arduino drives the `SER` input to the first shift register; the `h'` output from this chip drives the `SER` input to the next chip and so-on down the line.

The `h'` output from the last chip in the chain is left dangling (unconnected);

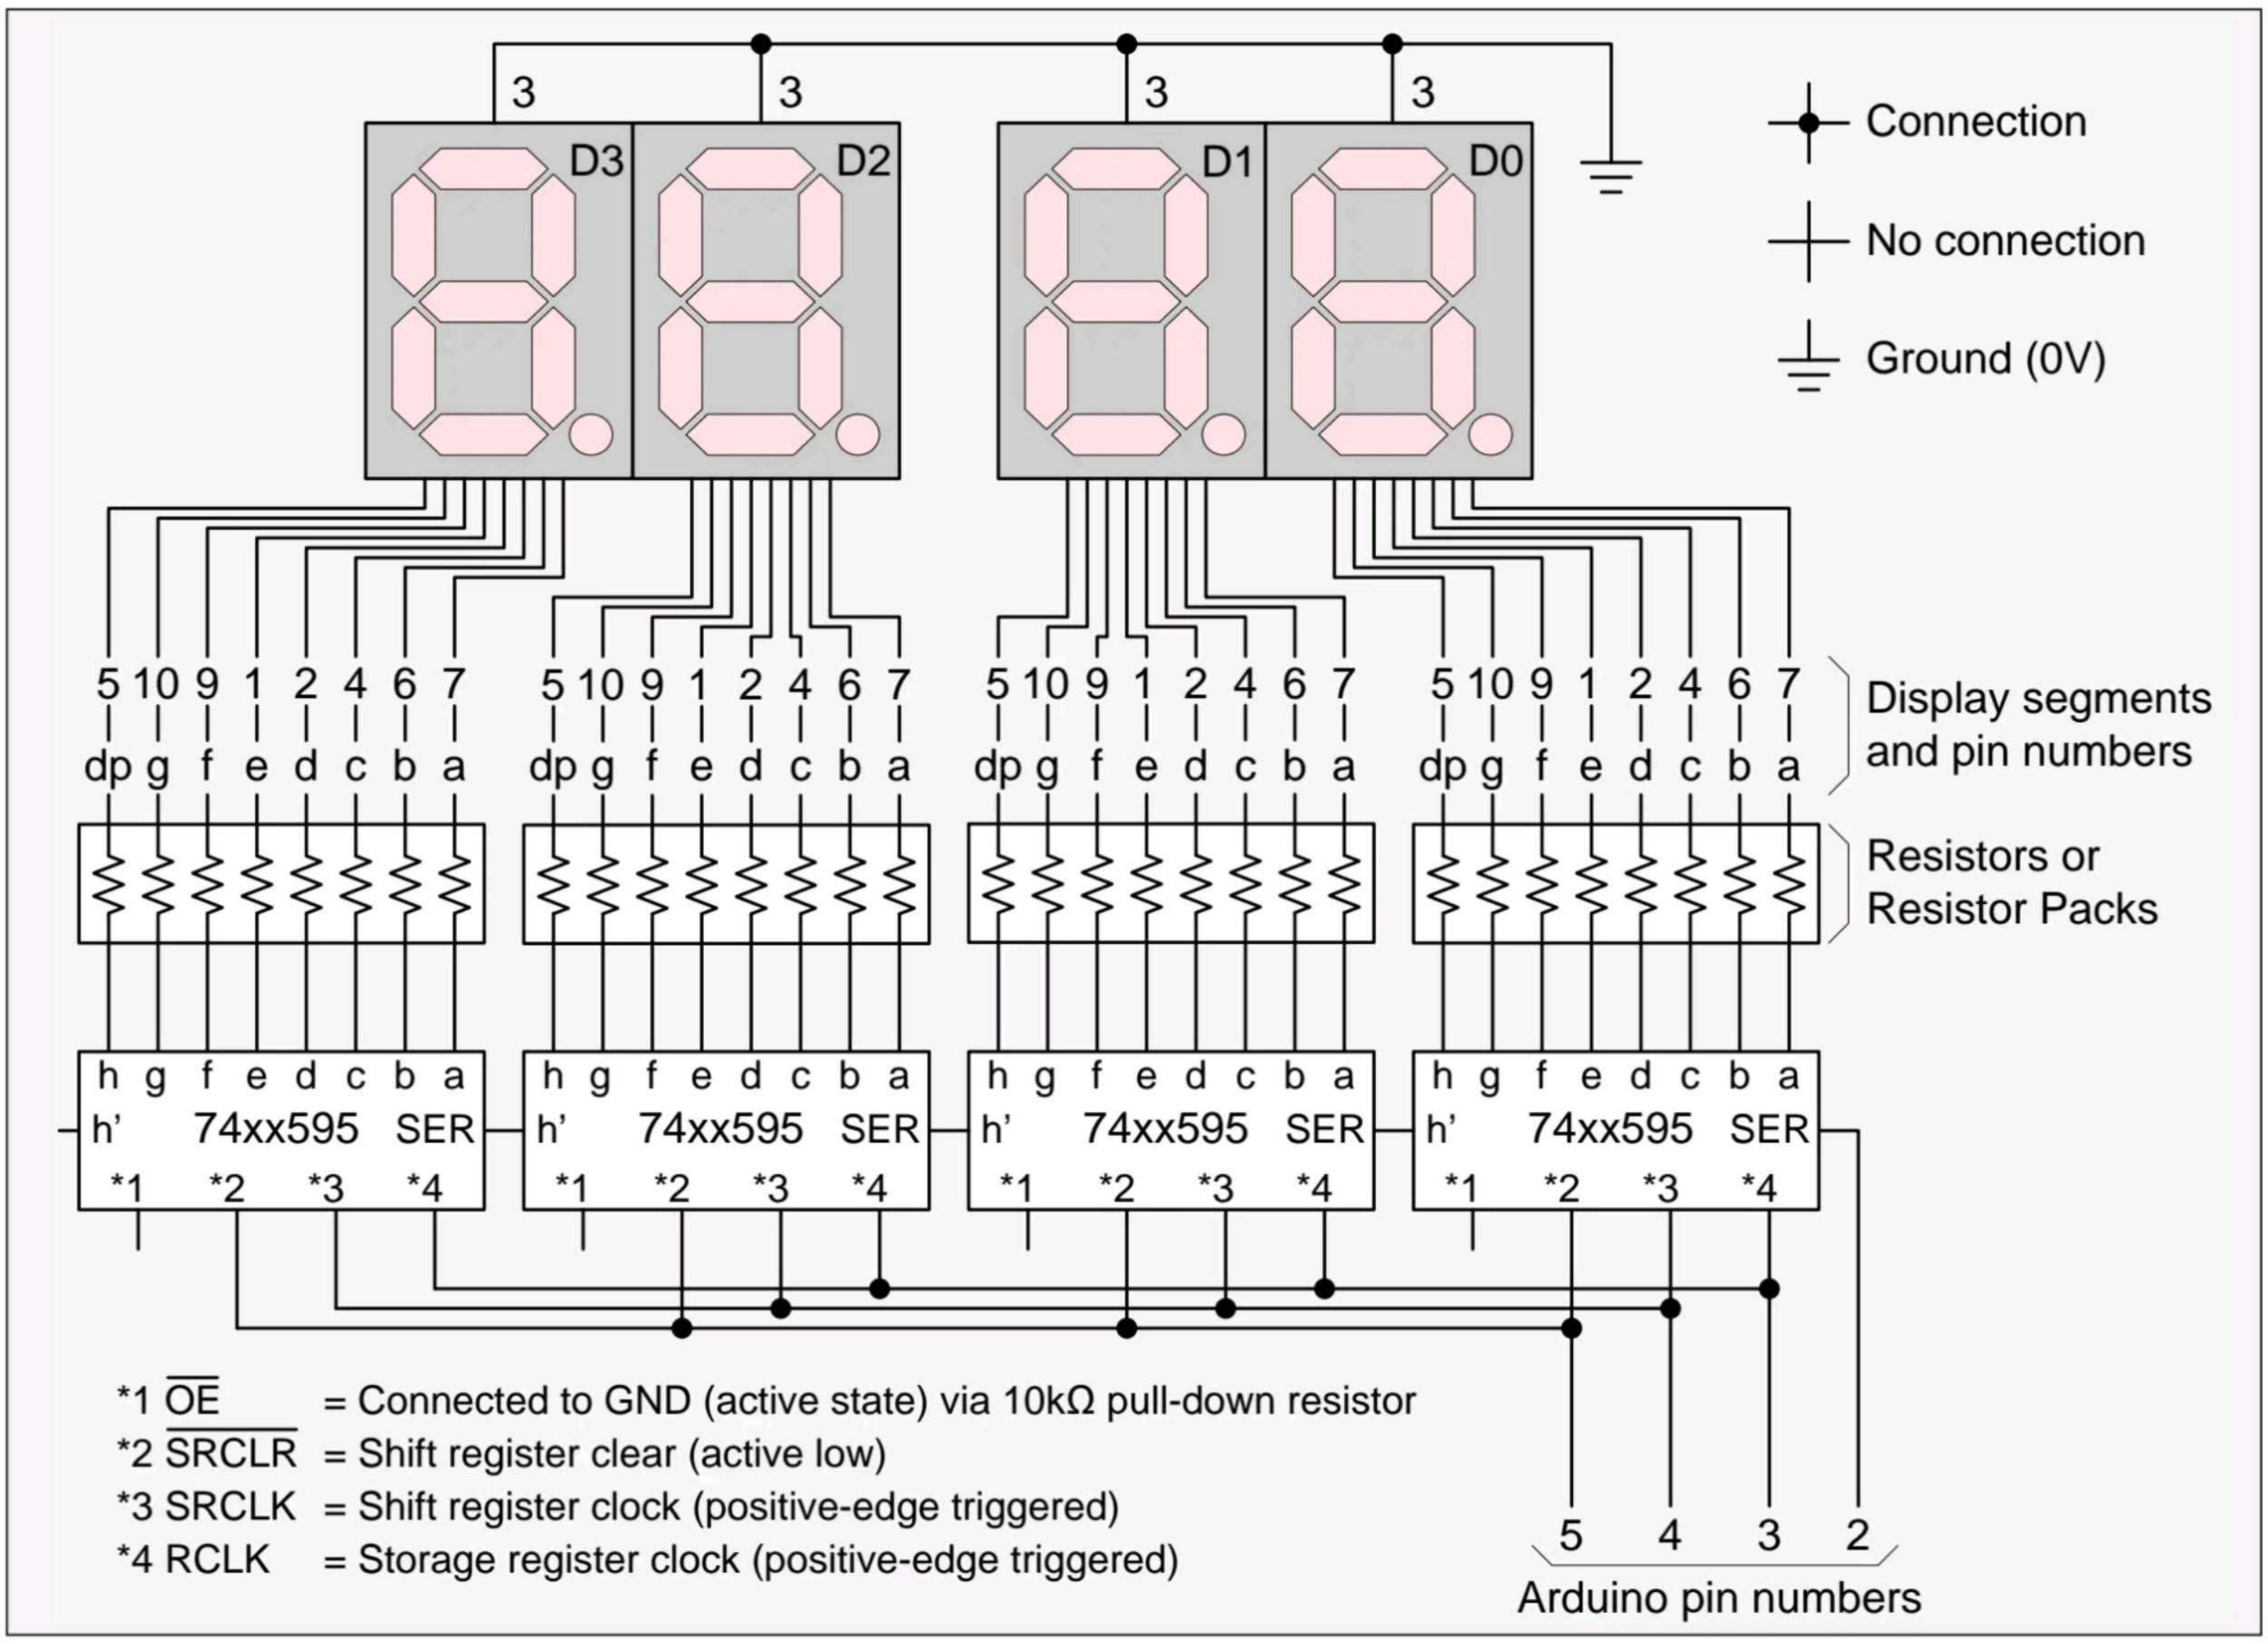
3
3
3
3
D3
D2
D1
D0
Connection
No connection
Ground (0V)
5 10 9 1 2 4 6 7
dp g f e d c b a
Display segments and pin numbers
Resistors or Resistor Packs
h g f e d c b a
h'
74xx595
SER
*1 *2 *3 *4
*1 $\overline{\text{OE}}$ = Connected to GND (active state) via 10kΩ pull-down resistor
*2 $\overline{\text{SRCLR}}$ = Shift register clear (active low)
*3 SRCLK = Shift register clock (positive-edge triggered)
*4 RCLK = Storage register clock (positive-edge triggered)
5 4 3 2
Arduino pin numbers

Fig.7: our new display driver circuit uses four 74xx595 shift registers.

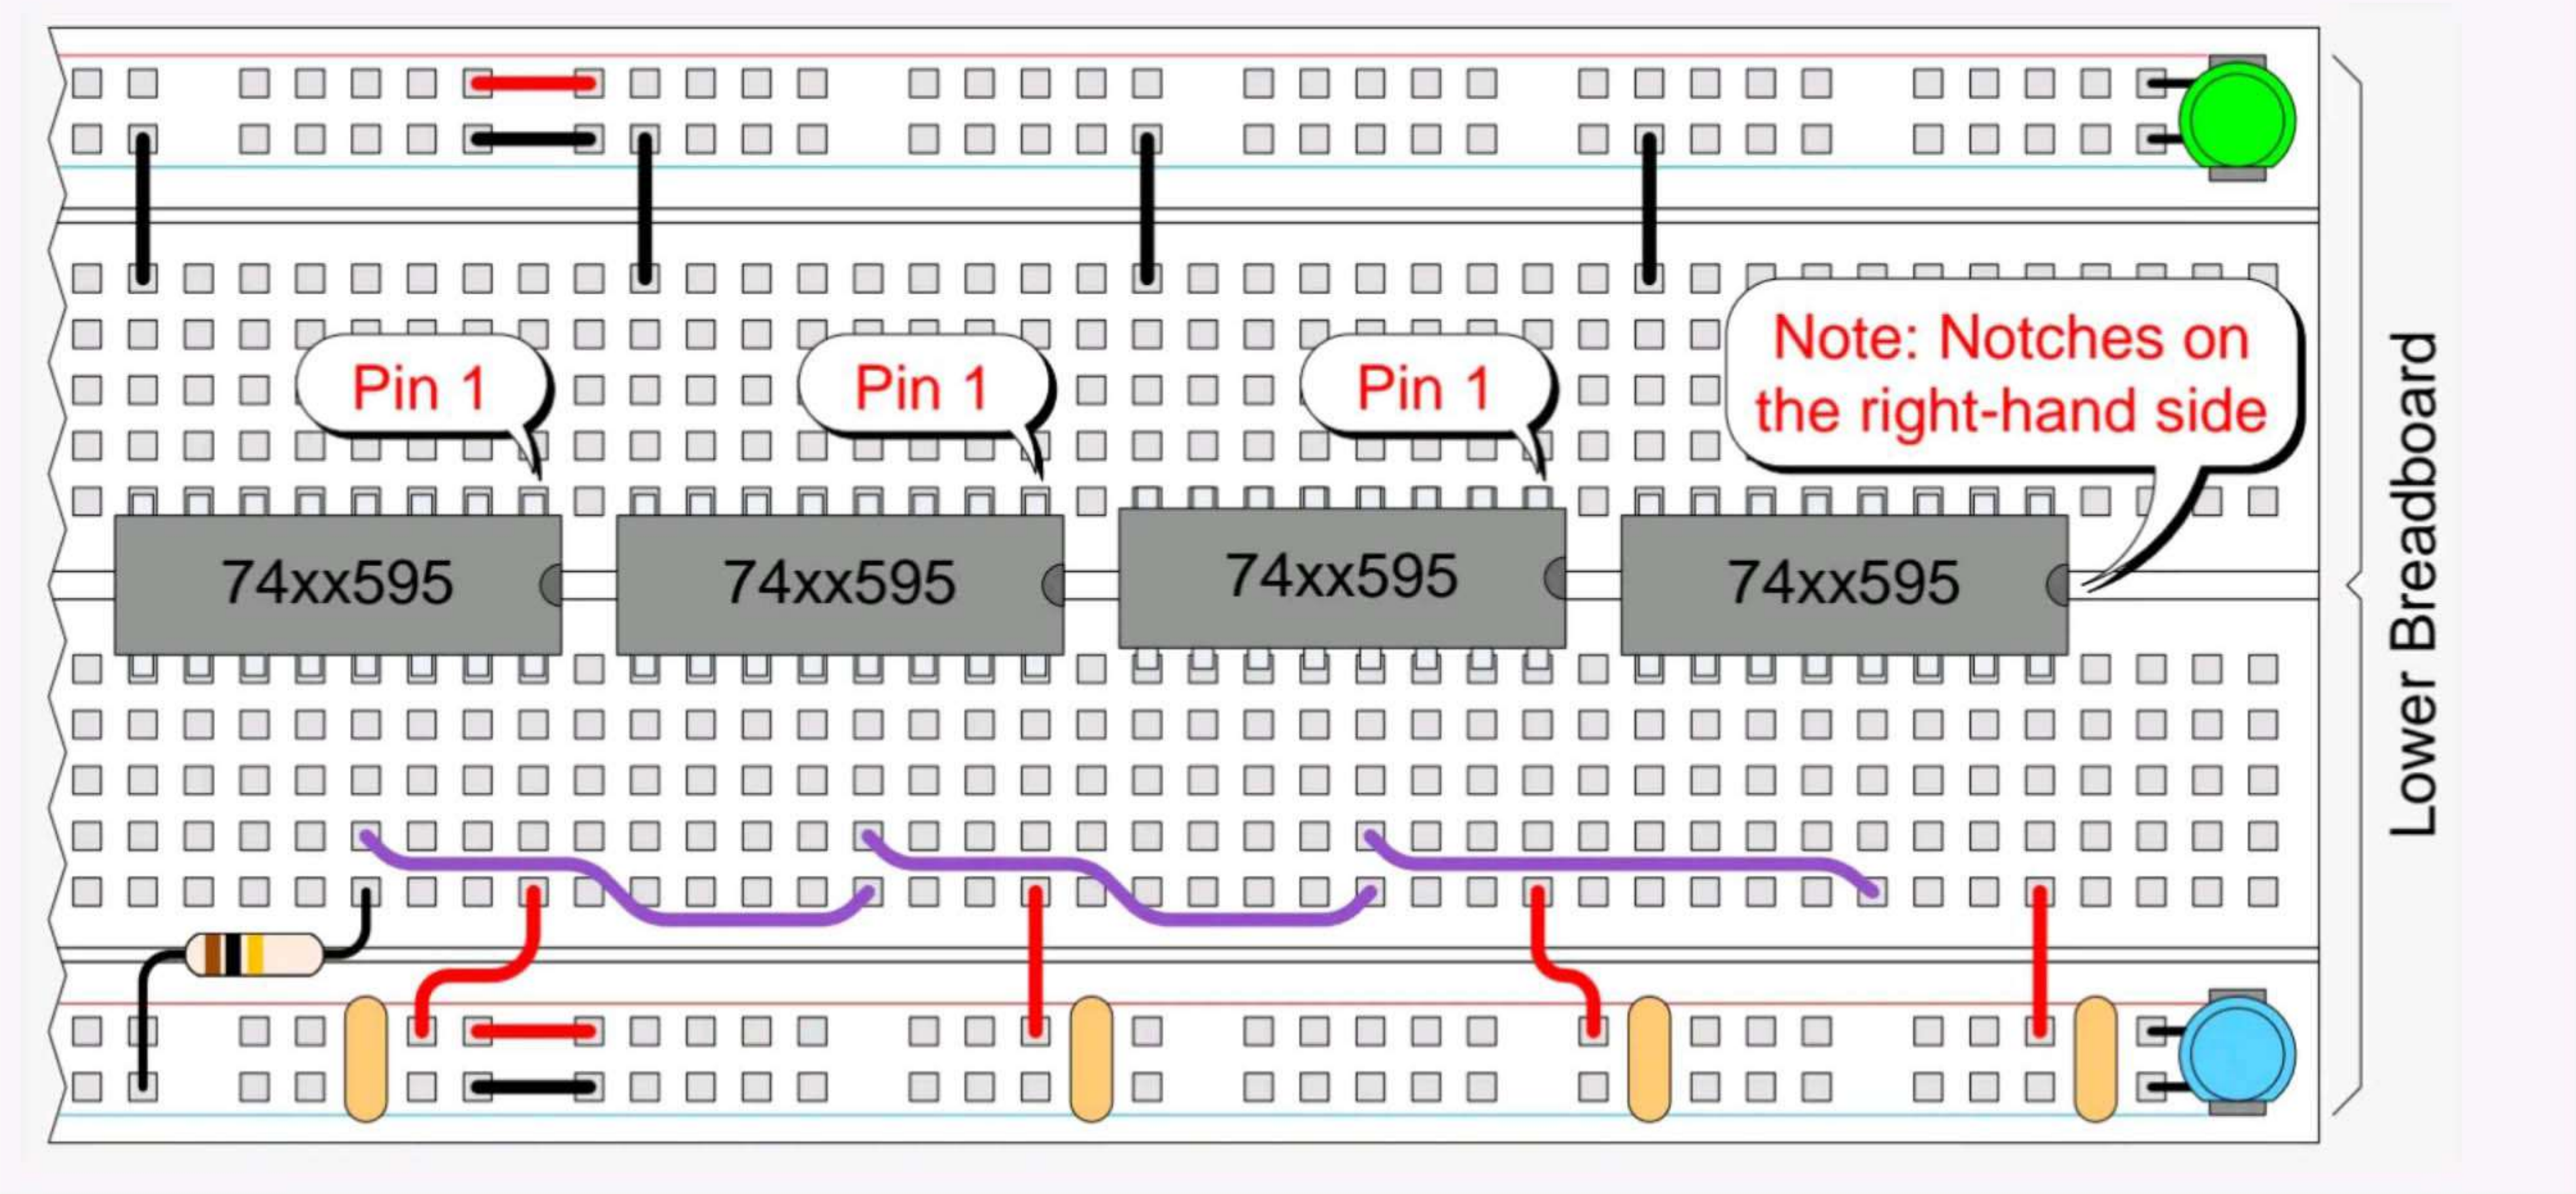

Fig.8: adding the shift registers.

we can use it to add more chips to the chain later if we wish.

The first thing we'll do is clear the decks for action. Make sure your clock is powered down, then remove the wires connecting Arduino pins 2 through 9 to our breadboards. We'll also remove any wires connected to the inputs of the CD4511 chips, plus any wires connecting the CD4511s to the four 270Ω resistor packs.

As part of this, we will remove any components (pull-up resistors and decoupling capacitors) associated with the CD4511s. Also, any power and ground wires associated with the CD4511s and the wires connecting pin 1 on each of the resistor packs to ground (associated with the `dp` segments on the displays).

Finally, we will remove the CD4511 chips and the resistor packs.

Resistance is futile

It would be wonderful if we could keep on using our existing 270Ω resistor packs, but it was not to be. When we look at the 74xx595 data sheet, we see that each output can supply only 6mA.

We know that the supply voltage from the Arduino Uno is 5V and that our red segment LEDs have a voltage drop of 2V. Rearranging Ohm's law, V = IR, giving R = V/I, tells us that our ideal resistors would be (5V – 2V) ÷ 0.006A = 500Ω.

The closest standard resistor value is 510Ω (with green-brown-brown color bands), which would be great, although 560Ω (with green-blue-brown color bands) would also do nicely.

Now we have a choice. We can use 32 discrete (individual) 510Ω or 560Ω resistors if we have them, either plugged directly into our breadboards, or on four 8-resistor breakout boards (BOBs) as discussed in *PE*, July 2024.

Another option would be to use 560Ω resistor packs, which is what I've shown in my diagrams.

As we've discussed before, these components are created by a variety of manufacturers, like the Bourns 4100R Series (**https://pemag.au/link/ac13**), and they are available from most major vendors, such as DigiKey (**https://pemag.au/link/ac1p**).

Adding the shift registers

Once we've added our new current-limiting resistors, the next step is to add the four shift registers as illustrated in Fig.8.

It's **very important** to note that we are rotating these chips 180° to the way we previously inserted our BCD decoders.

This is a tricky topic. The commonly accepted practice is to orientate our chips horizontally such that the notch/dimple is on the left and pin 1 is located on the lower left-hand corner. However, we are positioning our shift register chips the other way round, with the notch/dimple on the right and pin 1 on the upper right-hand corner.

We are doing this to make our wiring easier (but I still feel naughty).

Add the black wires connecting pin 8 on these chips to the 0V (ground) rail and the red wires connecting pin 16 to the +5V (power) rail. Also add the 100nF decoupling capacitors between the power and ground rails as close to the red wires as possible.

Add the 10kΩ pull-down resistor and the purple wires as shown. This ensures that pin 13 (the $\overline{\text{OE}}$ input) to each device is in its active low state, thereby enabling the outputs.

You can see the current state of play in the file **CB-Dec24-brd-02.pdf**

This would be a great time to power everything up and use your multimeter to verify that you are seeing +5V (give or take) across all the power and ground rail pairs on the breadboards. Also check that you are seeing +5V (give or take) between pins 8 and 16 on each chip (apply the black probe to pin 8 and the red probe to pin 16).

Keeping the red probe on pin 16, also check that you are seeing +5V between pin 16 and pin 13 on each chip. Once you are satisfied that everything is tickety-boo, power the system down again before proceeding to the next step.

Adding the wires

The next step is to add the wires driving the inputs to the shift registers as illustrated in Fig.9. Although these wires are all shown as being orange in these diagrams, I personally prefer to use different colors in my own implementation.

For example, I started by using a long yellow wire to connect pin 2 on the Arduino to pin 14 (the `SER` data input) on the right-hand shift register. I then used a short yellow wire to connect pin 9 (the $Q_{H'}$ output) on the right-hand shift register to pin 14 (the `SER` input) on the next chip in the chain.

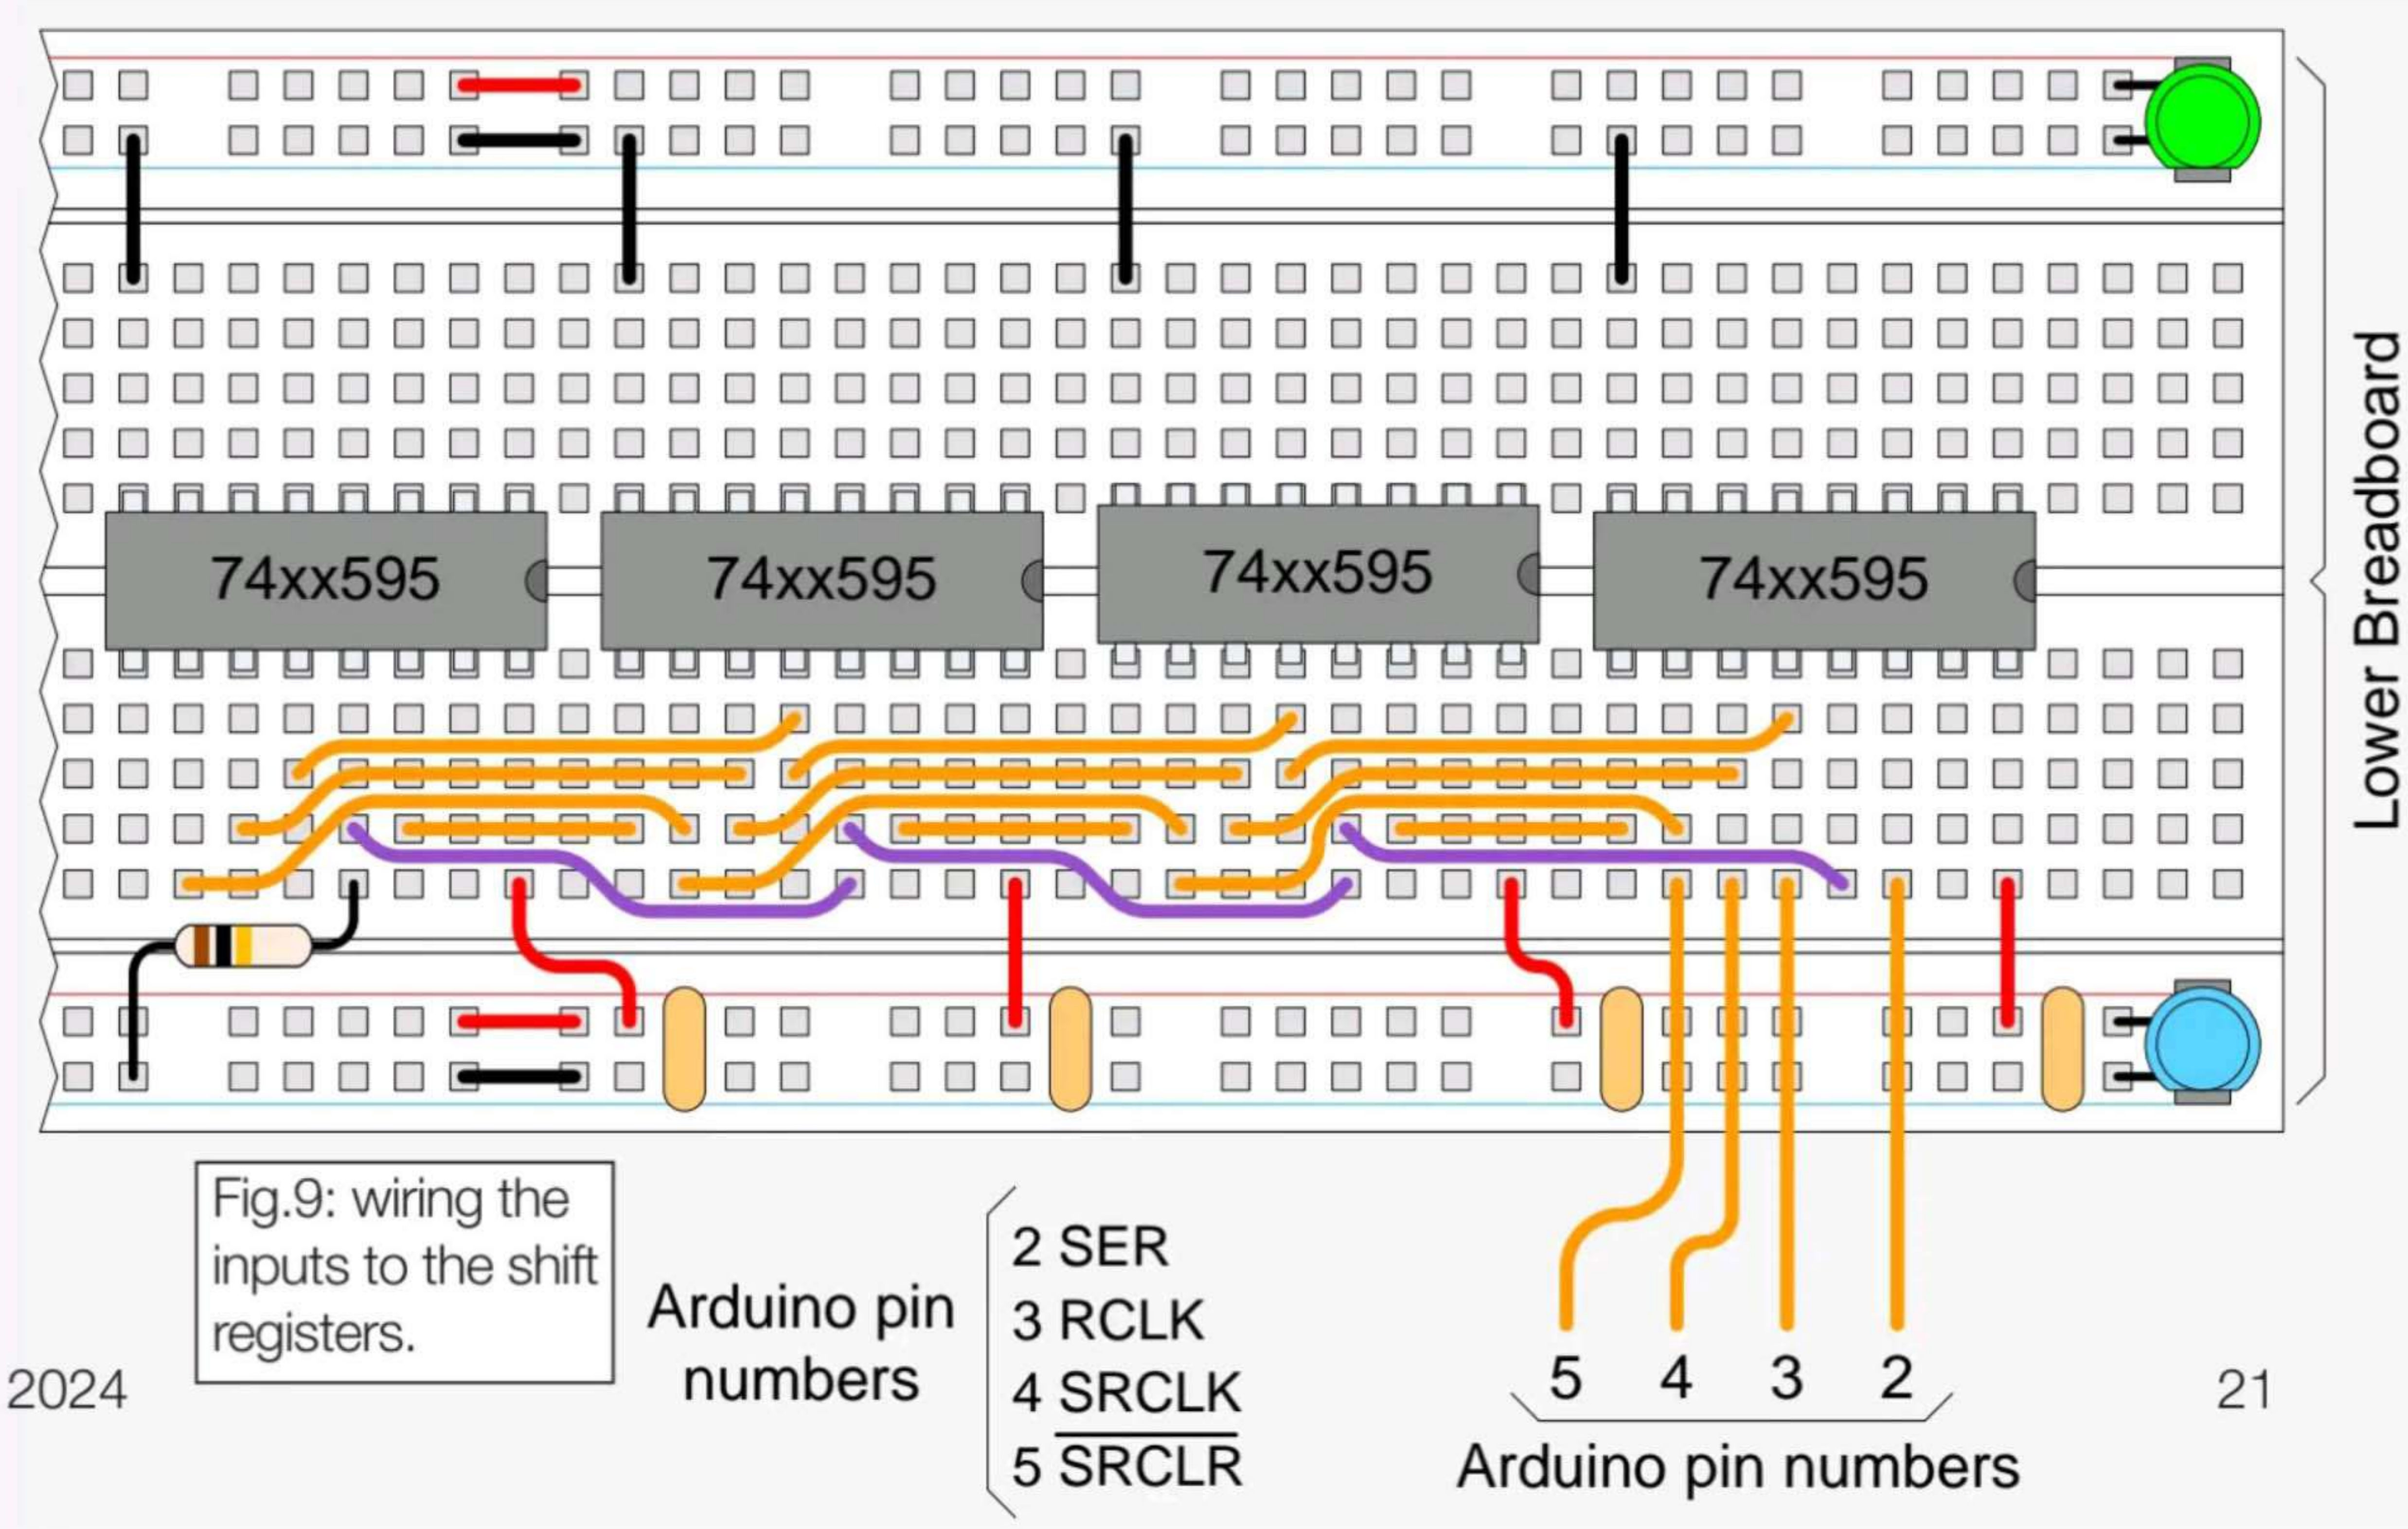

Fig.9: wiring the inputs to the shift registers.

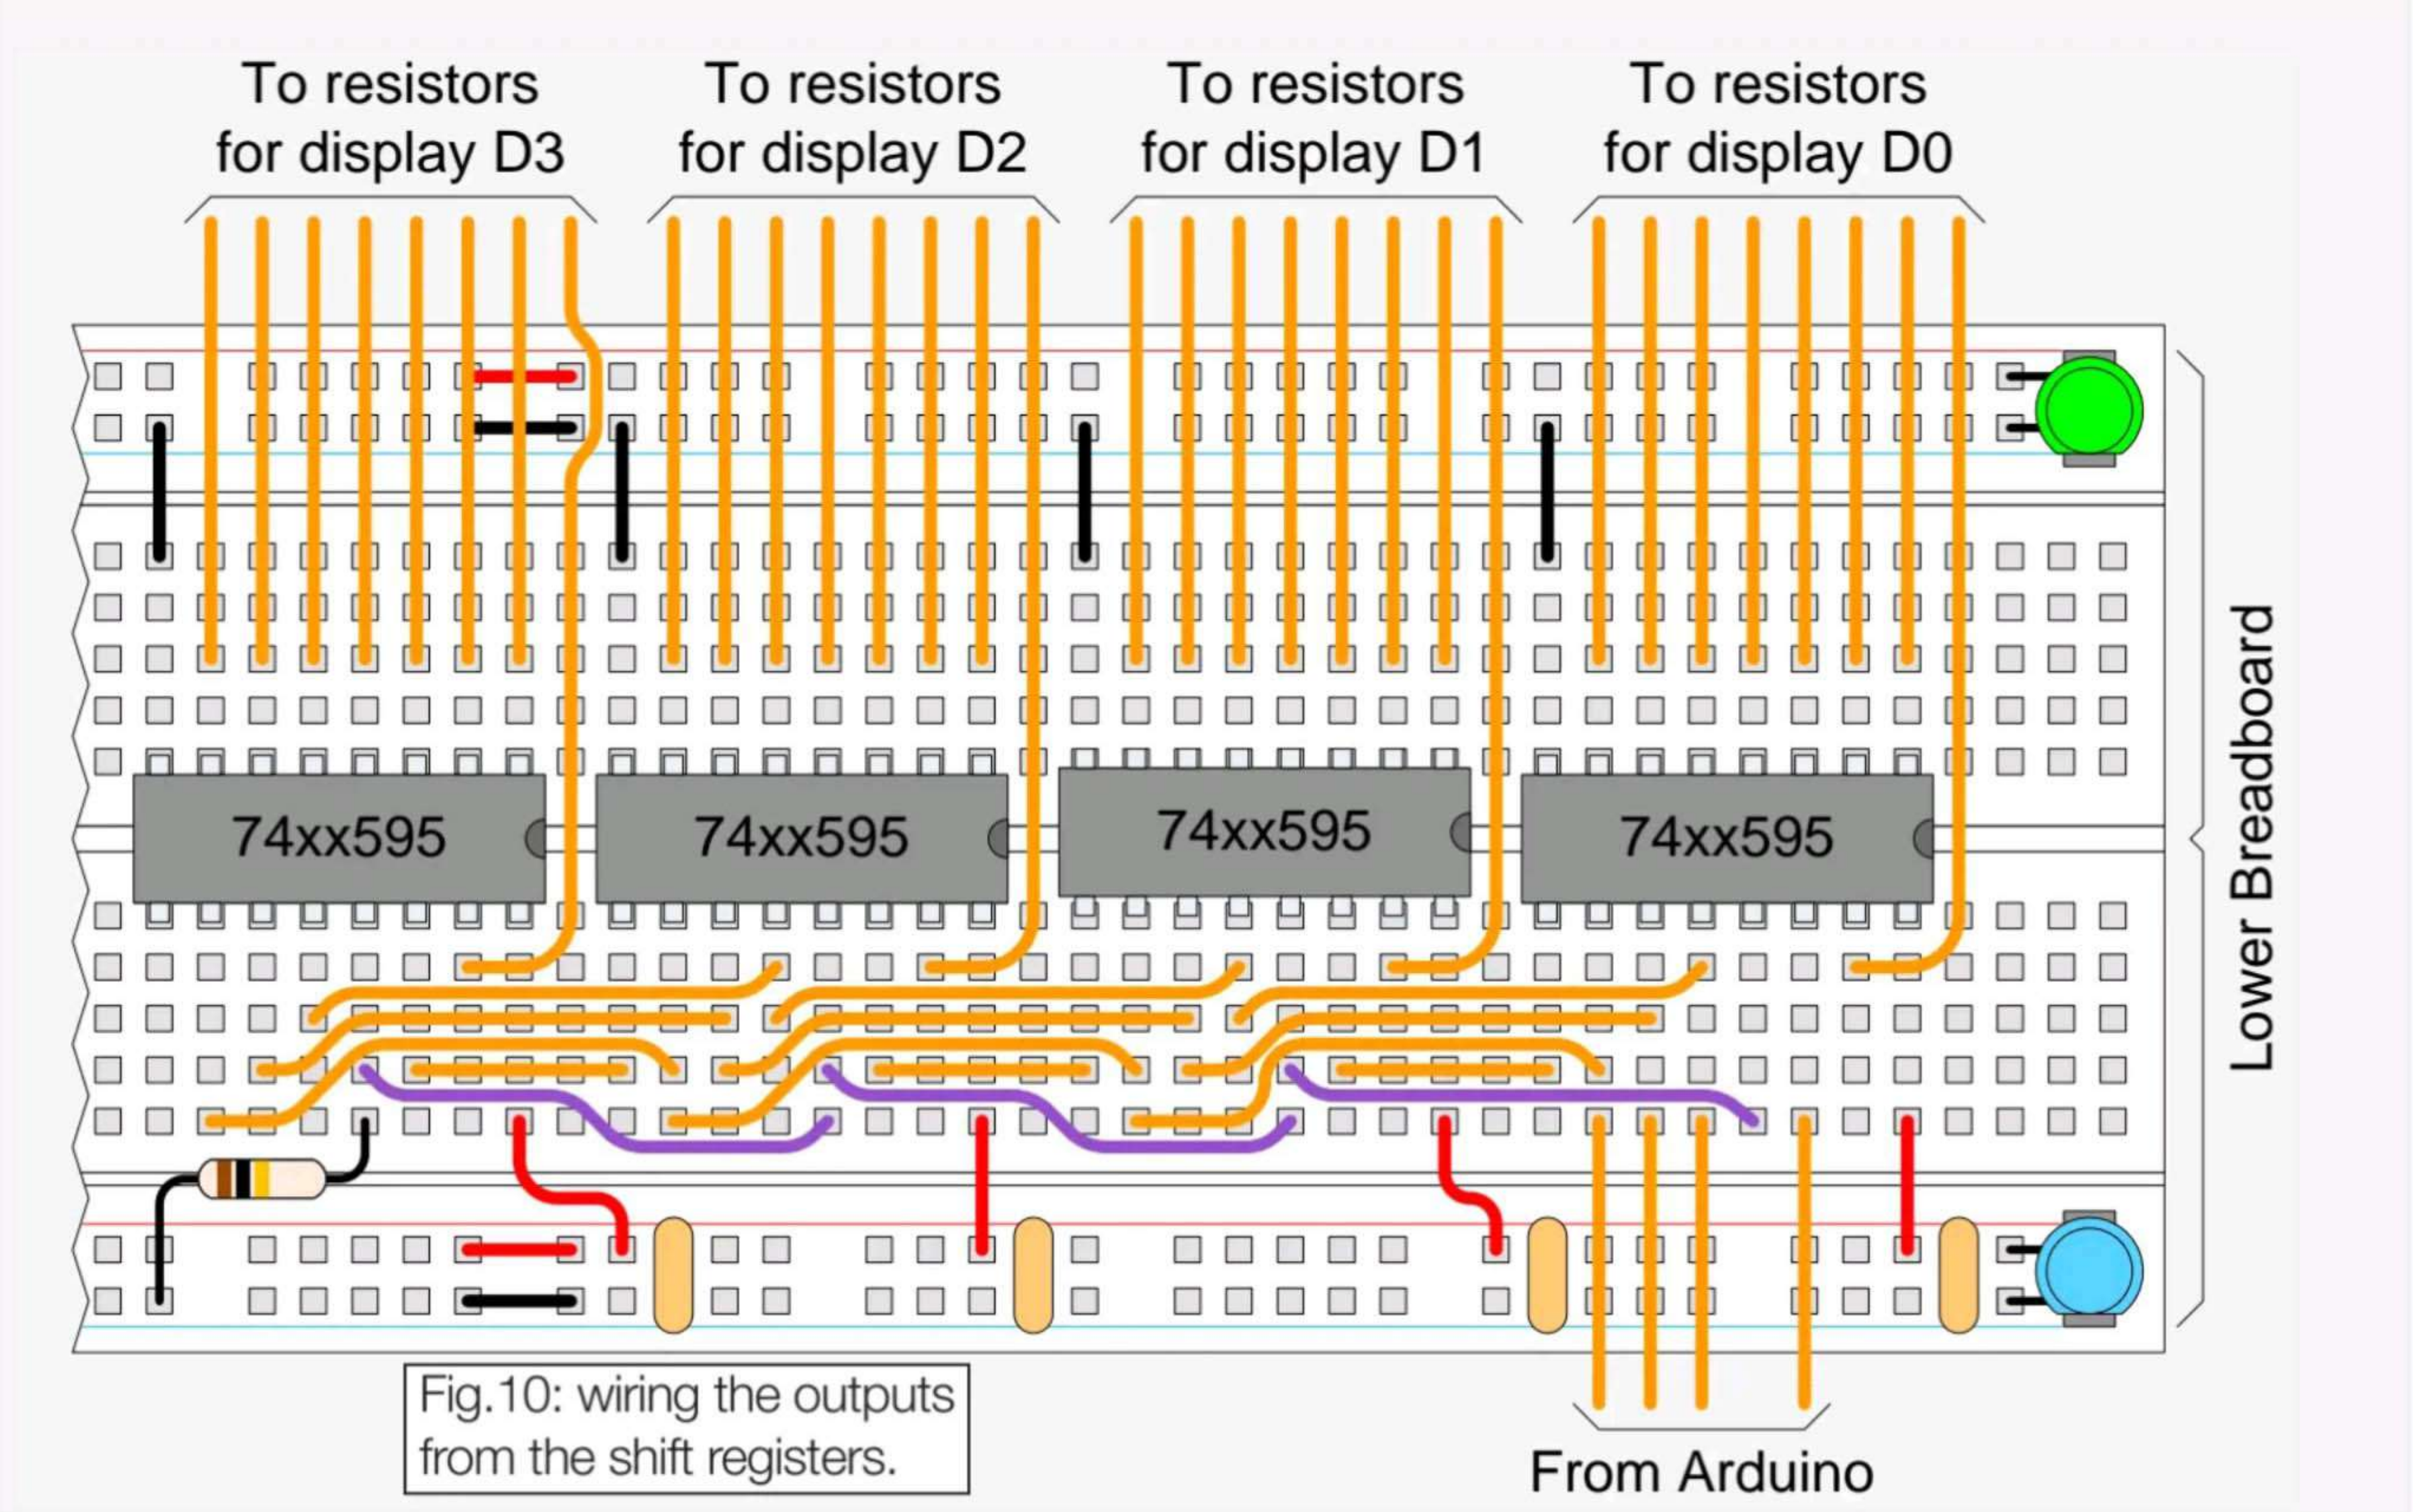
To resistors for display D3
To resistors for display D2
To resistors for display D1
To resistors for display D0
74xx595
74xx595
74xx595
74xx595
Lower Breadboard
From Arduino

Fig.10: wiring the outputs from the shift registers.

Similarly, I used two more short yellow wires to connect the $Q_{H'}$ output from the second chip to the `SER` input on the third chip, and the $Q_{H'}$ output from the third chip to the `SER` input on the fourth.

Next, I connected pin 3 from the Arduino to pin 12 (the `RCLK` input) on the right-hand chip, then added short wires to connect this pin to the corresponding pins on the other chips. Similarly, I connected the `SRCLK` and $\overline{\texttt{SRCLR}}$ signals as shown.

We're so close I can taste it. The final step is to connect the outputs from the shift registers to the inputs to the current-limiting resistors driving the displays (Fig.10). The full-up implementation is shown in file **CB-Dec24-brd-03.pdf**

A quick test

Before we 'go for broke', let's perform a quick test to ensure we can talk to our shift register chain in a language it understands. I whipped up a little test program, which you can find in the file **CB-dec24-code-01.txt**

As you'll see when you cast your orbs over this code, we start with some definitions that will ease our lives by making our program more understandable and easier to modify in the future.

In addition to defining `NUM_SEGS_ALL` to be 32 (the total number of segments in our four displays), we also define the active and inactive values for our `RCLK`, `SRCLK`, and $\overline{\texttt{SRCLR}}$ signals.

Next, we declare the names and numbers of the Arduino pins we wish to use to drive our shift register chain, assigning pins 2, 3, 4, and 5 to variable names `PinSer`, `PinRclk`, `PinSrclk` and `PinSrclr`, respectively.

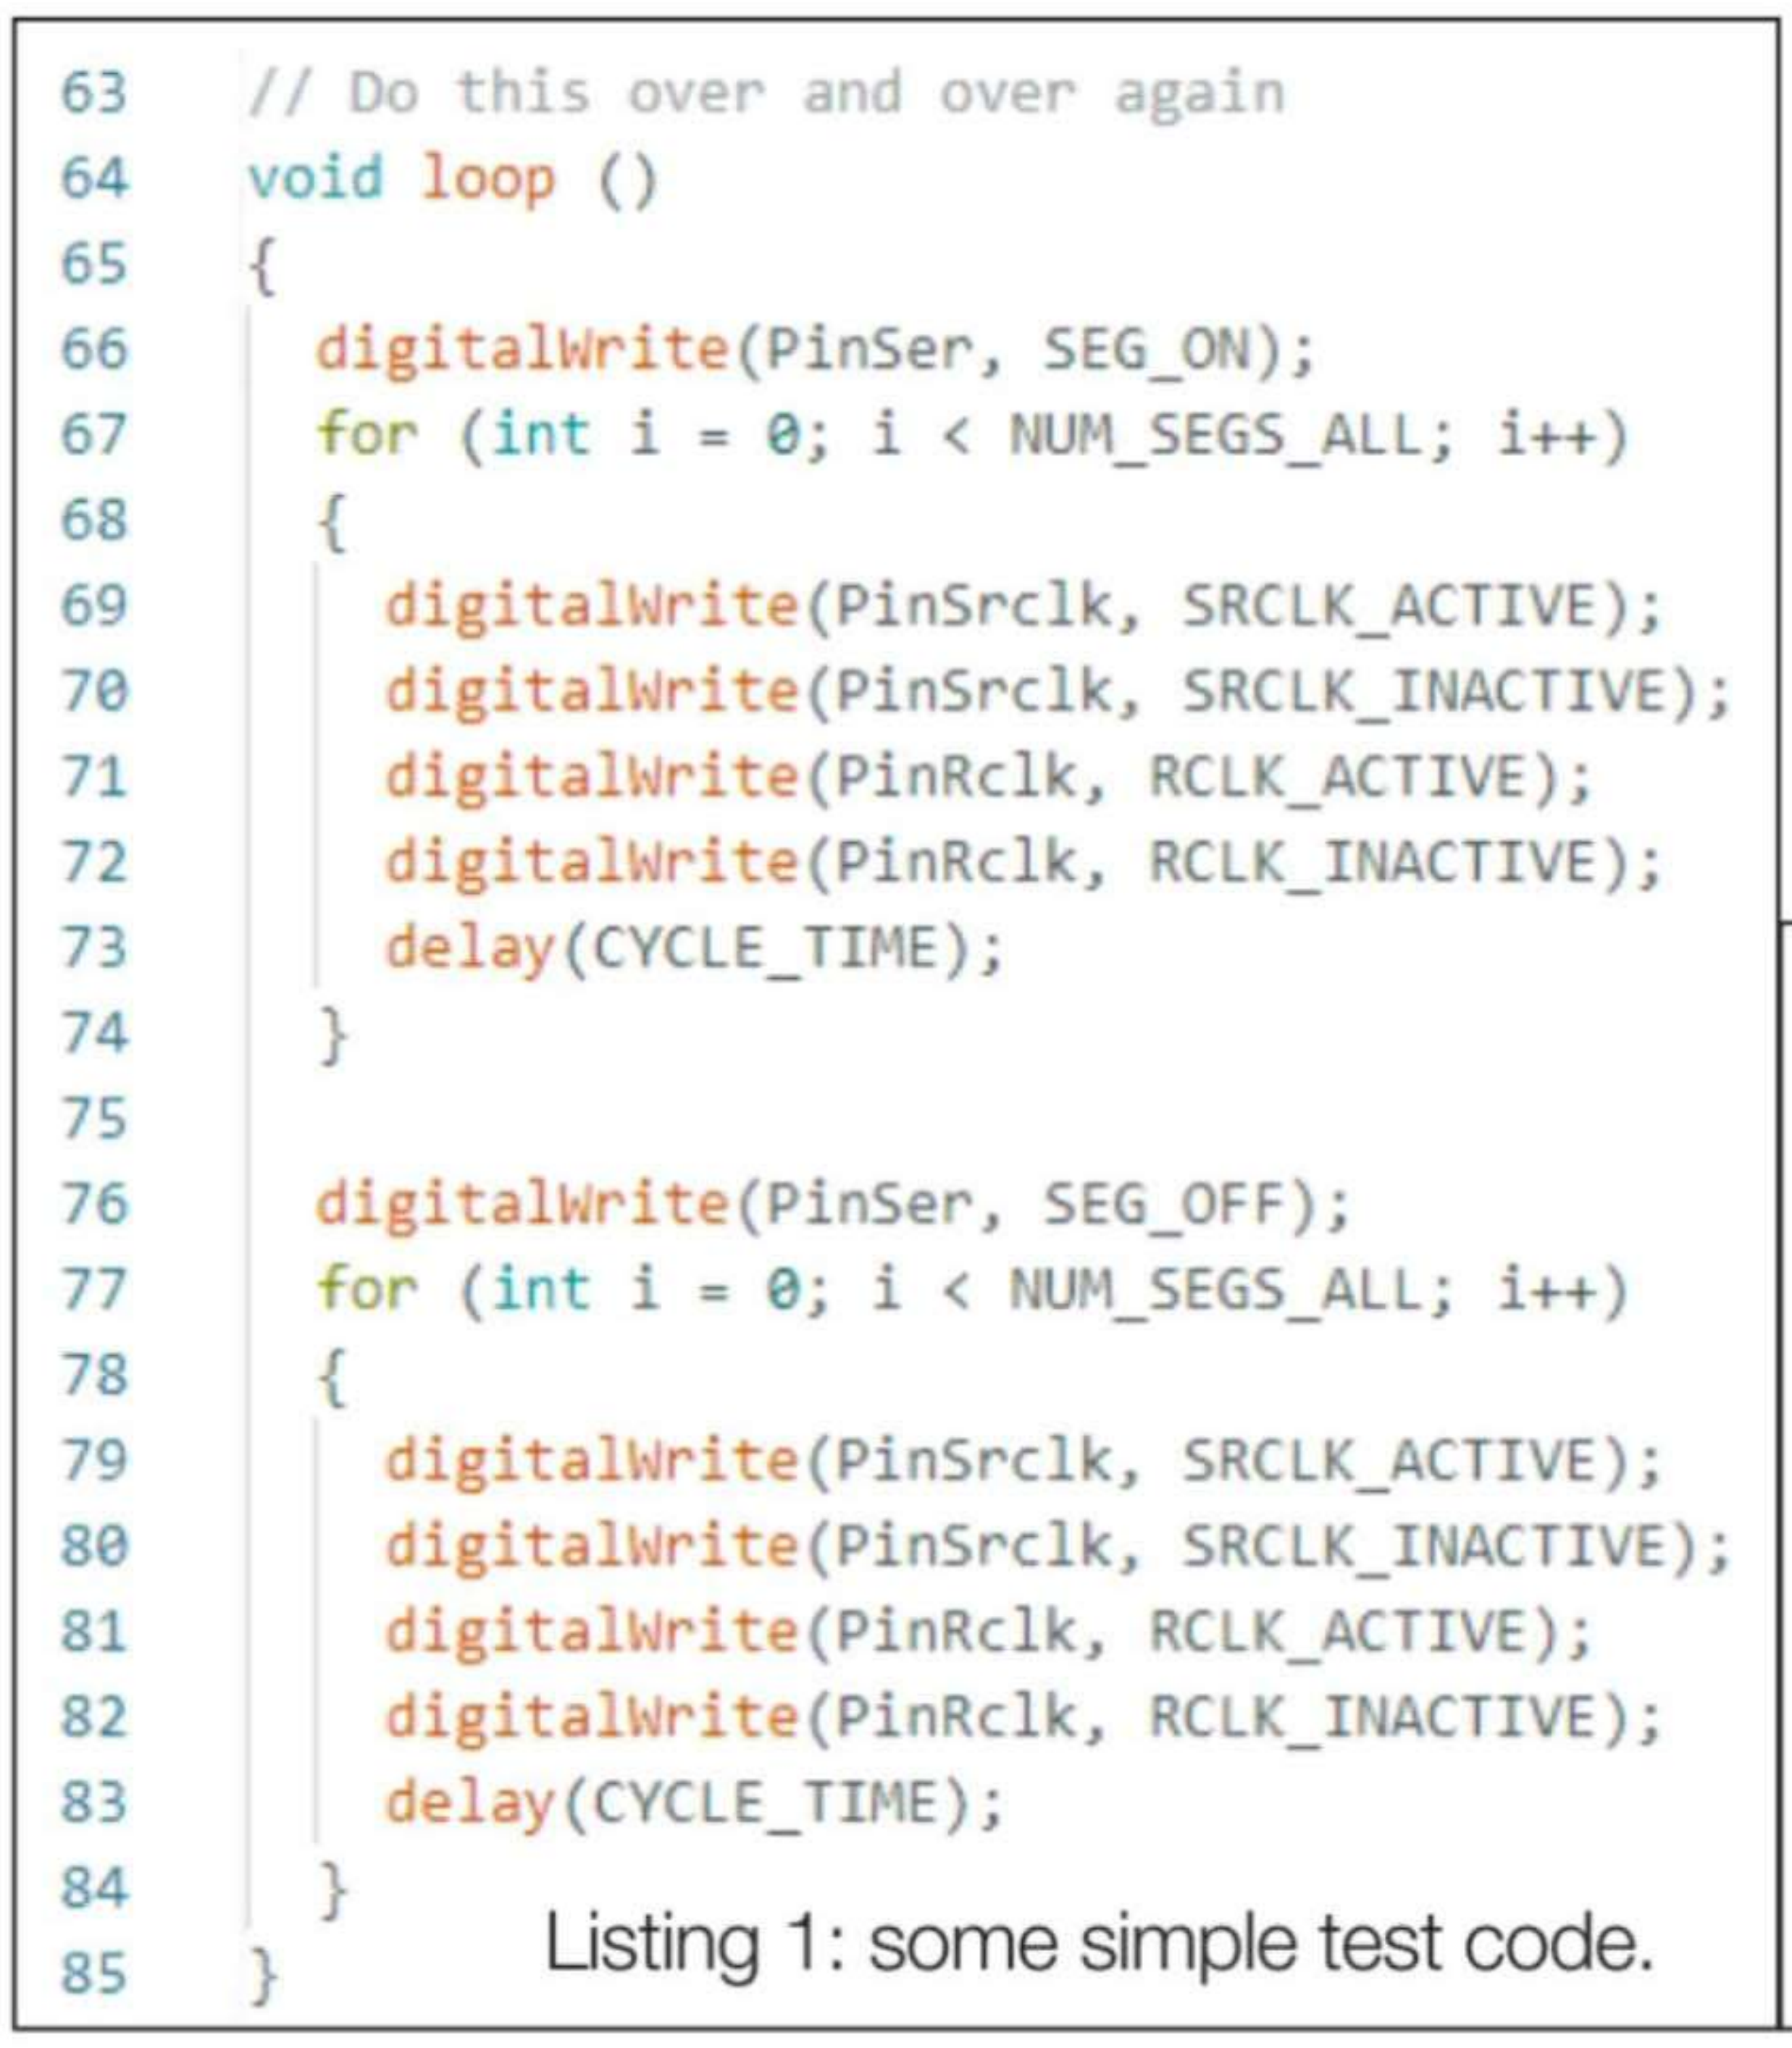

```
63  // Do this over and over again
64  void loop ()
65  {
66    digitalWrite(PinSer, SEG_ON);
67    for (int i = 0; i < NUM_SEGS_ALL; i++)
68    {
69      digitalWrite(PinSrclk, SRCLK_ACTIVE);
70      digitalWrite(PinSrclk, SRCLK_INACTIVE);
71      digitalWrite(PinRclk, RCLK_ACTIVE);
72      digitalWrite(PinRclk, RCLK_INACTIVE);
73      delay(CYCLE_TIME);
74    }
75
76    digitalWrite(PinSer, SEG_OFF);
77    for (int i = 0; i < NUM_SEGS_ALL; i++)
78    {
79      digitalWrite(PinSrclk, SRCLK_ACTIVE);
80      digitalWrite(PinSrclk, SRCLK_INACTIVE);
81      digitalWrite(PinRclk, RCLK_ACTIVE);
82      digitalWrite(PinRclk, RCLK_INACTIVE);
83      delay(CYCLE_TIME);
84    }
85  }
```

Listing 1: some simple test code.

In our `setup()` function, we first set all the shift register signals to their inactive values, then we pulse the `PinSrclr` signal to clear the shift register chips, after which we pulse the `PinRclk` signal to copy the contents of the shift registers into their output registers.

The portion of the code of most interest to us here is the `loop()` function, shown in Listing 1. We start on line 66 when we present a logic 1 to the serial input to the first shift register in the chain. Then we use the `for()` loop on lines 67 through 74 to repeatedly load this value into the shift register.

Inside this loop, we first pulse the `PinSrclk` signal to copy the data value into the shift register, then we pulse the `PinRclk` pin to copy the contents of the shift registers into their output registers. We wait for a second to give ourselves time to see what's happening on our displays, then we do it all again.

We expect to see the segments on the displays light up one after the other, starting with `a`, then `b`, then `c`… ending with `dp` on the first display, followed by the other displays in their turn.

On line 76, we present a logic 0 to the serial input to the first shift register in the chain, then we use the `for()` loop on lines 77 through 84 to turn all the segments off again, one after the other.

Once you've verified that everything is working as expected, why don't you experiment with this program a little. For example, if you delete lines 71 through 73 from inside the first loop and reinsert them outside the loop starting at line 75, all the segments will light up simultaneously, stay lit for one second, then start to be extinguished one after the other as before.

Cunning code

Now take a deep breath because the time has come to pull everything together. We're going to take our old clock program, which was based on our BCD decoder implementation, and modify it to use our new shift register realisation. You can feast your orbs on my version of this code in the file **CB-dec24-code-02.txt**

After all we've been through together, you should be able to understand the bulk of this program without any problem. The interesting part is found in the `loop()` function.

We start by declaring a four-element array of 8-bit unsigned integers called `tmpData[]` (temporary data), along with a 32-bit unsigned integer called `dispData` (display data). Observe that we are using the `uint8_t` and `uint32_t` data types for these values (we introduced these data types in detail in *PE*, September 2020).

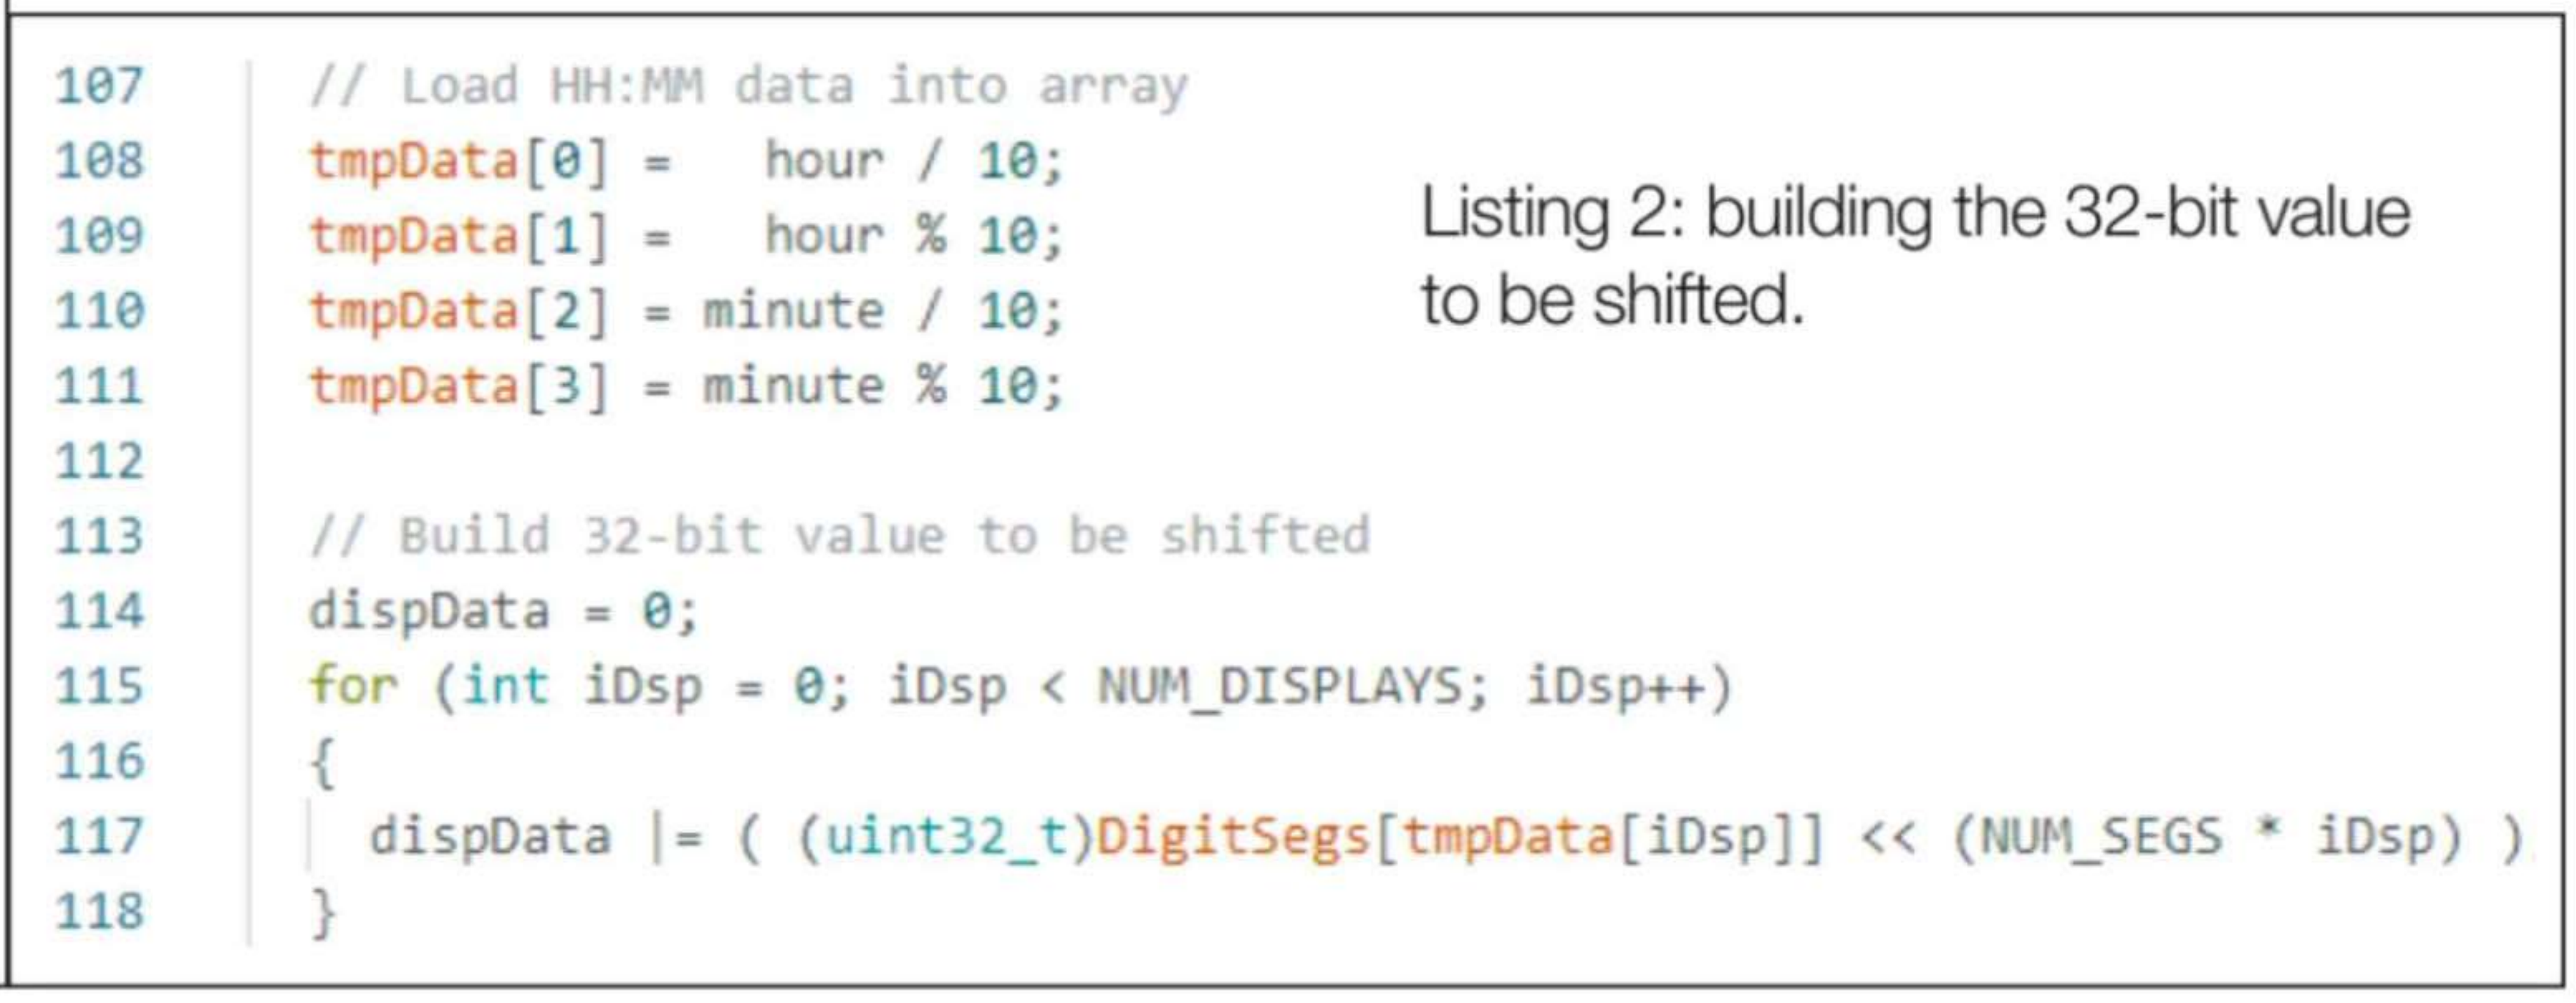

```
107  // Load HH:MM data into array
108  tmpData[0] =   hour / 10;
109  tmpData[1] =   hour % 10;
110  tmpData[2] = minute / 10;
111  tmpData[3] = minute % 10;
112
113  // Build 32-bit value to be shifted
114  dispData = 0;
115  for (int iDsp = 0; iDsp < NUM_DISPLAYS; iDsp++)
116  {
117    dispData |= ( (uint32_t)DigitSegs[tmpData[iDsp]] << (NUM_SEGS * iDsp) )
118  }
```

Listing 2: building the 32-bit value to be shifted.

As usual, we use our real-time clock (RTC) to determine the hour (in 24-hour format) and minute, both of which are represented as integers. As usual, we convert it from 24-hour format to 12-hour format and we load the most- and least-significant digits of our hour and minute values into our temporary data array as illustrated in lines 108 through 111 of Listing 2.

Now what we need to do is build the 32-bit value to be loaded into our shift register. Remember earlier when I wrote that we need to load the last bit first and the first bit last? That explains why we load the most-significant hour digit into `tmpData[0]` and the least-significant minute digit into `tmpData[3]`, opposite to the way we did things before.

Our next task is to take these four 8-bit `tmpData[]` values and use them to construct a single 32-bit `dispData` value, as illustrated in Fig.11.

We start on line 114 by clearing our `dispData` value to all 0s. The next bit is quite clever, because we use a `for()` loop whose index follows the sequence 0, 1, 2, 3. Each time round the loop, we use this index to shift the contents of the corresponding `tmpData[]` value left by the index multiple of 8 bits using the << shift left operator.

This means `tmpData[0]` will be shifted left by 8*0 = 0 bits, `tmpData[1]` will be shifted left by 8*1 = 8 bits, `tmpData[2]` will be shifted left by 8*2 = 16 bits and so forth. As part of this, 0s will be shifted into the newly-vacated, least-significant (right-hand) bits.

After performing each shift, we logically OR the existing contents of `dispData` with the contents of the newly shifted value using the compound |= assignment operator. If you wish to delve deeper into any of this, you might take a look at my *Masking and the C/C++ Bitwise Operators* column (https://pemag.au/link/ac1q).

One part of this deserves special attention. Observe the use of `(uint32_t)` on line 117. Let's consider what would happen if we omitted this from our statement.

"Type casting" refers to the process of converting one data type to another. There are two forms of type casting: implicit and explicit. One example of an implicit type cast occurs when we assign a smaller value to a larger one, like assigning an 8-bit value to a 32-bit value, for example. In this case, the 8-bit value will be padded with 24 leading zeros (assuming it's an unsigned type, anyway).

In line 117 of Listing 2, this implicit cast is the last thing that occurs prior to the new value being assigned to `dispData`. Before this, our 8-bit `tmpData[]` values will be shifted left by multiples of 8 bits as we just discussed.

Since the << shift left operator has a higher precedence than the |= assignment operator, the shifted bits will 'fall off the end' of our 8-bit value.

This is why we are using `(uint32_t)` to perform an explicit cast. It has a higher precedence than a shift operator, so our 8-bit `tmpData[]` values will be cast into 32-bit values (padded with 0s in the most-significant bits) before the shift takes place. It's the resulting shifted 32-bit value that will be assigned to our 32-bit `dispData` variable.

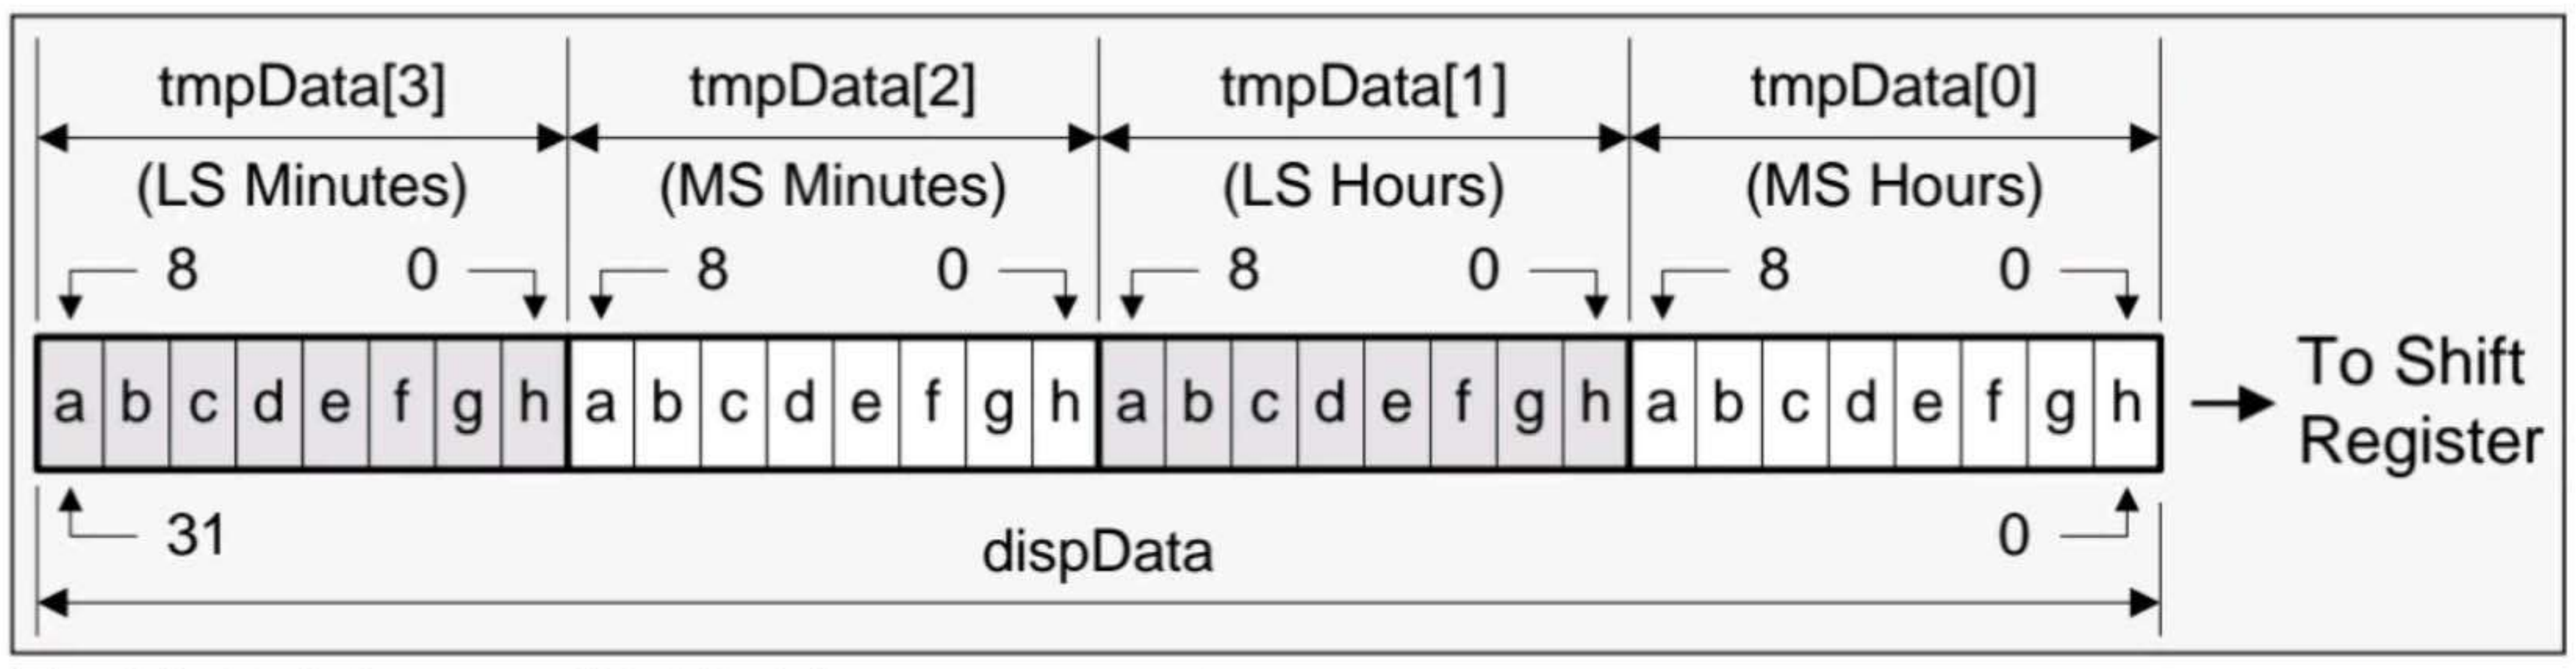

Fig.11: Building our 32-bit `dispData` value.

Components from Part 1	
LEDs (assorted colours)	https://amzn.to/3E7VAQE
Resistors (assorted values)	https://amzn.to/3O4RvBt
Solderless breadboard	https://amzn.to/3O2L3e8
Multicore jumper wires (male-to-male)	https://amzn.to/3O4hnxk
Components from Part 2	
7-segment display(s)	https://amzn.to/3Afm8yu
Components from Part 5	
Momentary pushbutton switches	https://amzn.to/3Tk7Q87
Components from Part 6	
Passive piezoelectric buzzer	https://amzn.to/3KmxjcX
Components for Part 9	
SW-18010P vibration switch	https://bit.ly/46SfDA4
Components for Part 10	
Breadboard mounting trimpots	https://bit.ly/3QAuz04
Components for Part 12	
Light-Dependent Resistor	https://bit.ly/3S2430m
Components for Part 13	
BC337 NPN Transistor	https://bit.ly/40xAgyS
Components for Part 14	
HC-SR04 Ultrasonic Sensor	https://bit.ly/49AMBq4
Components for Part 15	
Real-Time Clock (RTC)	https://bit.ly/3S9OjHl
Components for Part 18	
Long tailed (0.1-inch/2.54mm pitch) header pins	https://bit.ly/3U1Vp2z
Components for Part 19	
Prototyping boards	https://bit.ly/3UMkcZ1
Kit of popular SN74LS00 chips	https://bit.ly/3wqgzyv
Components for Part 20	
16V 100µF electrolytic capacitors	https://bit.ly/44LzpNa
Ceramic capacitors (assorted values)	https://bit.ly/4bEAUiv
Components for Part 22	
SN74LS48N BCD Decoders	https://bit.ly/3zT18jx
16-pin resistor pack (8 × 390Ω)	https://bit.ly/4d0ISDz
Components for Part 23	
4 × CD4511 BCD Decoders	https://bit.ly/3yzWl6k
16-pin resistor pack (8 × 390Ω)	https://bit.ly/46HZ2PQ
Components for Part 24	
74HC595 8-bit shift registers	https://bit.ly/3yzWl6k
4 × 16-pin resistor packs (8 × 560Ω)	https://bit.ly/46HZ2PQ

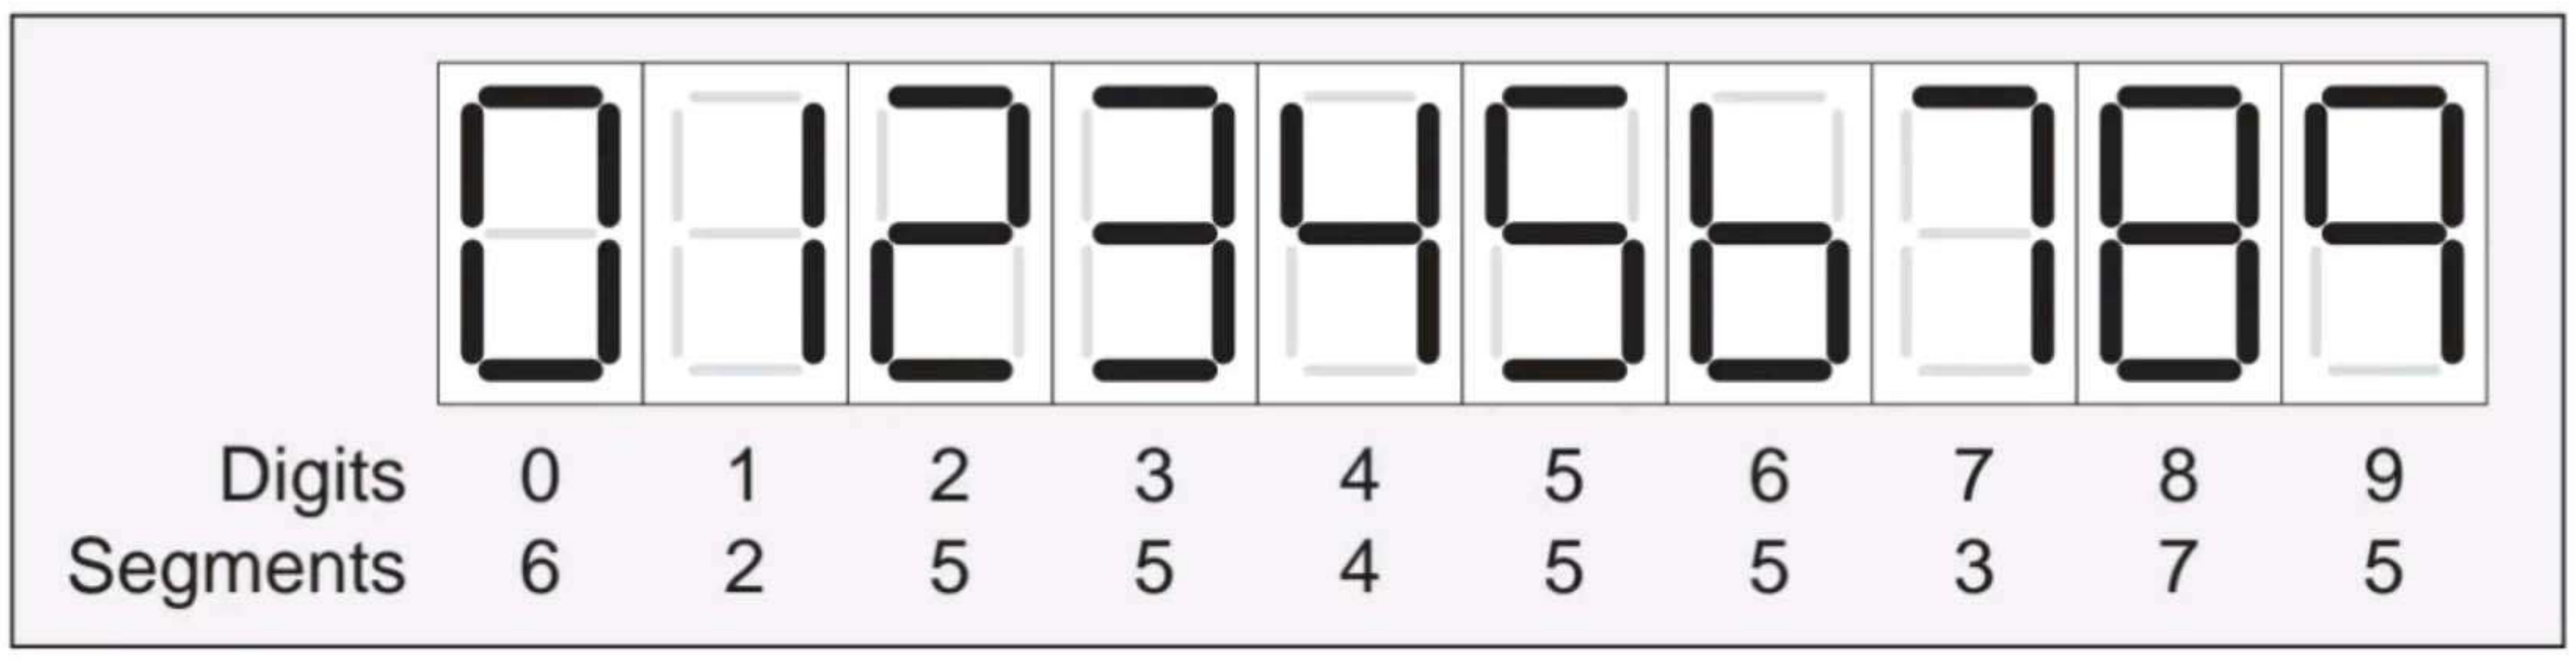

Digits	0	1	2	3	4	5	6	7	8	9
Segments	6	2	5	5	4	5	5	3	7	5

Fig.12: the number of segments used by each digit in our current display scheme.

The last function in my program is `DisplayTime()`, which takes our 32-bit `dispData` value and loads it into the shift register, starting with bit 0 and ending with bit 31. That means bit 0 in `dispData` ends up driving the `dp` segment in display D3, while bit 31 in `dispData` ends up driving the `a` segment in display D0.

I just powered up my clock, and it works like a charm, as is apparent from the great big beaming smile on my face (I look like this now: 🙂).

All is revealed!

Now we consider our concluding conundrum. Our first task was to determine the maximum number of segments we might see active at the same time. Before we start, write your solution down on a piece of paper (we don't want any cheating).

Some people's minds work in such a way that the answer will simply pop out to them. Others, like your humble narrator, are obliged to take a more circuitous approach. There is no right or wrong path to success here. I started by noting the number of segments associated with each displayed digit (Fig.12).

Based on this, I created the tables shown in Fig.13. The reason we don't need to document all the possibilities in the minutes table is that the most-significant digit cycles through 0, 1, 2, 3, 4, 5. Since we know 0 consumes the most segments out of these options, we can remove the remaining combinations (10 through 59) from further consideration.

As we see, our worst-case scenario occurs at eight minutes past eight o'clock in the evening, which is 20:08 when our clock is running in its 24-hour mode. This will cause us to have 11 + 13 = 24 segments lit at the same time.

I could have solved this in my head, but would I have been sure I was right? By jotting things down in this way, we can see how we arrived at our answer and easily check things later.

Our second task was to determine the value of the current-limiting resistors that will provide us with the maximum possible segment current value assuming a total allocation of 350mA (the origin of this value was discussed in our previous column).

We want to drive our segments as bright as possible without exceeding their 20mA maximum as defined in their data sheet. I'm afraid to say that this is a trick question because the answer depends on the type of chips we are using to drive our displays.

Let's start by calculating 350mA ÷ 24 segments = 14.58mA per segment. If we were still using the CD4511 BCD to 7-segment decoder chips from our previous column, each output can handle up to 25mA, which means our 14.58mA will be a breeze.

Since we are using the Arduino Uno's 5V power supply, and we know the forward voltage drop of our red LED segments is 2V, we can calculate the desired value of our current-limiting resistors as (5V – 2V) ÷ 0.01458A = 206Ω. The nearest regular resistor value is 200Ω (with red-black-brown color bands), which is "close enough for government work", as they say, so we would go with that.

However, since we are now using 74HC595 shift registers to drive our displays, we are bound to obey the constraints of **their** data sheet, which states that the maximum current per output is only 6mA.

That still results in relatively bright displays, and it reduces our worst-case display current to 6mA × 24 segments = 144mA. This is well below our 350mA maximum, leaving 350mA – 144mA = 206mA on the table for us to use for other purposes.

However, it also means that the values of our current-limiting resistors are now calculated as (5V – 2V) ÷ 0.006A = 500Ω. In this case, the nearest regular value is 510Ω (with green-brown-brown color bands).

What? That's it?

I'm afraid we've come to the end of our *Arduino Bootcamp* series. Still, turn that frown upside-down into a smile because all is not doom and gloom.

In this series, we've covered a lot of ground and learned all sorts of interesting things, many of which will prove useful in our new *Weird and Wonderful Arduino Projects* series, starting next month.

Until then, let's all think about things we could do to extend the capabilities of the clock project we just created. I'll share my thoughts in our next column (albeit under a different name).

As always, I welcome your insightful comments, penetrating questions, and sagacious suggestions. PE

12-Hour	
H	#Segs
1	2
2	5
3	5
4	4
5	5
6	5
7	3
8	7
9	5
10	2+6 = 8
11	2+2 = 4
12	2+5 = 7

24-Hour			
H	#Segs	H	#Segs
0	6	12	2+5 = 7
1	2	13	2+5 = 7
2	5	14	2+4 = 6
3	5	15	2+5 = 7
4	4	16	2+5 = 7
5	5	17	2+3 = 5
6	5	18	2+7 = 9
7	3	19	2+5 = 7
8	7	20	5+6 = 11
9	5	21	5+2 = 7
10	2+6 = 8	22	5+5 = 10
11	2+2 = 4	23	5+5 = 10

Minutes	
M	#Segs
00	6+6 = 12
01	6+2 = 8
02	6+5 = 11
03	6+5 = 11
04	6+4 = 10
05	6+5 = 11
06	6+5 = 11
07	6+3 = 9
08	6+7 = 13
09	6+5 = 11
10	2+6 = 8
:	: :
59	5+5 = 10

Fig.13: my approach to solving the problem.

Online resources

For the purposes of this series, I'm going to assume that you are already familiar with fundamental concepts like voltage, current and resistance. If not, you might want to start by perusing and pondering a short series of articles I penned on these very topics – see: https://pemag.au/link/ac16

Similarly, I'll assume you are no stranger to solderless breadboards. Having said this, even if you've used these little scamps before, there are some aspects to them that can trap the unwary, so may I suggest you feast your orbs on a column I wrote just for you – see: https://pemag.au/link/ac17

Last, but not least, you will find a treasure trove of resources at the **Arduino.cc** website, including example programs and reference documentation.

Exclusive offer

Win a Microchip dsPIC33C Digital Power Starter Kit

Practical Electronics is offering its readers the chance to win a Microchip dsPIC33C Digital Power Starter Kit (DM330017-3). Even if you don't win, receive a 15%-off voucher, plus free shipping for one of these products.

The dsPIC33C Digital Power Starter Kit demonstrates the capabilities and features of Microchip's SMPS device families.

Modern power supplies are becoming smaller, more efficient, more flexible and less expensive. These enhancements have come about as Digital Signal Controllers (DSCs) are incorporated into Switch Mode Power Supply (SMPS) designs.

The board provided in the kit is intended to introduce the features of the Microchip SMPS families of devices. The dsPIC33C Digital Power Starter Kit features an on-board dsPIC33CK256MP505 digital signal controller (DSC), SMPS power stages, resistive loads, LCD screen and USB/UART bridge. It also includes a programmer/debugger, eliminating the need for any additional hardware.

Worth US$291 (about £223) each

These single-core dsPIC33CK DSCs accelerate Digital Signal Processing (DSP) performance for time-critical applications. This family offers fast, deterministic performance to address real-world design requirements. Complementing the dual-core dsPIC33CH DSCs, the dsPIC33CK family offers a cost-effective single-core option with the same high-performance core and peripherals.

These high-performance DSCs feature a dsPIC33 'C' core with DSP engine, expanded context-selected registers to reduce interrupt latency, new instructions to accelerate DSP performance, tightly coupled peripherals and faster execution to enable complex high-speed control loop implementation. The dsPIC33CK family maintains the simplicity of a microcontroller and provides a performance upgrade path for dsPIC33E and dsPIC33F DSC users.

Free-to-enter competition

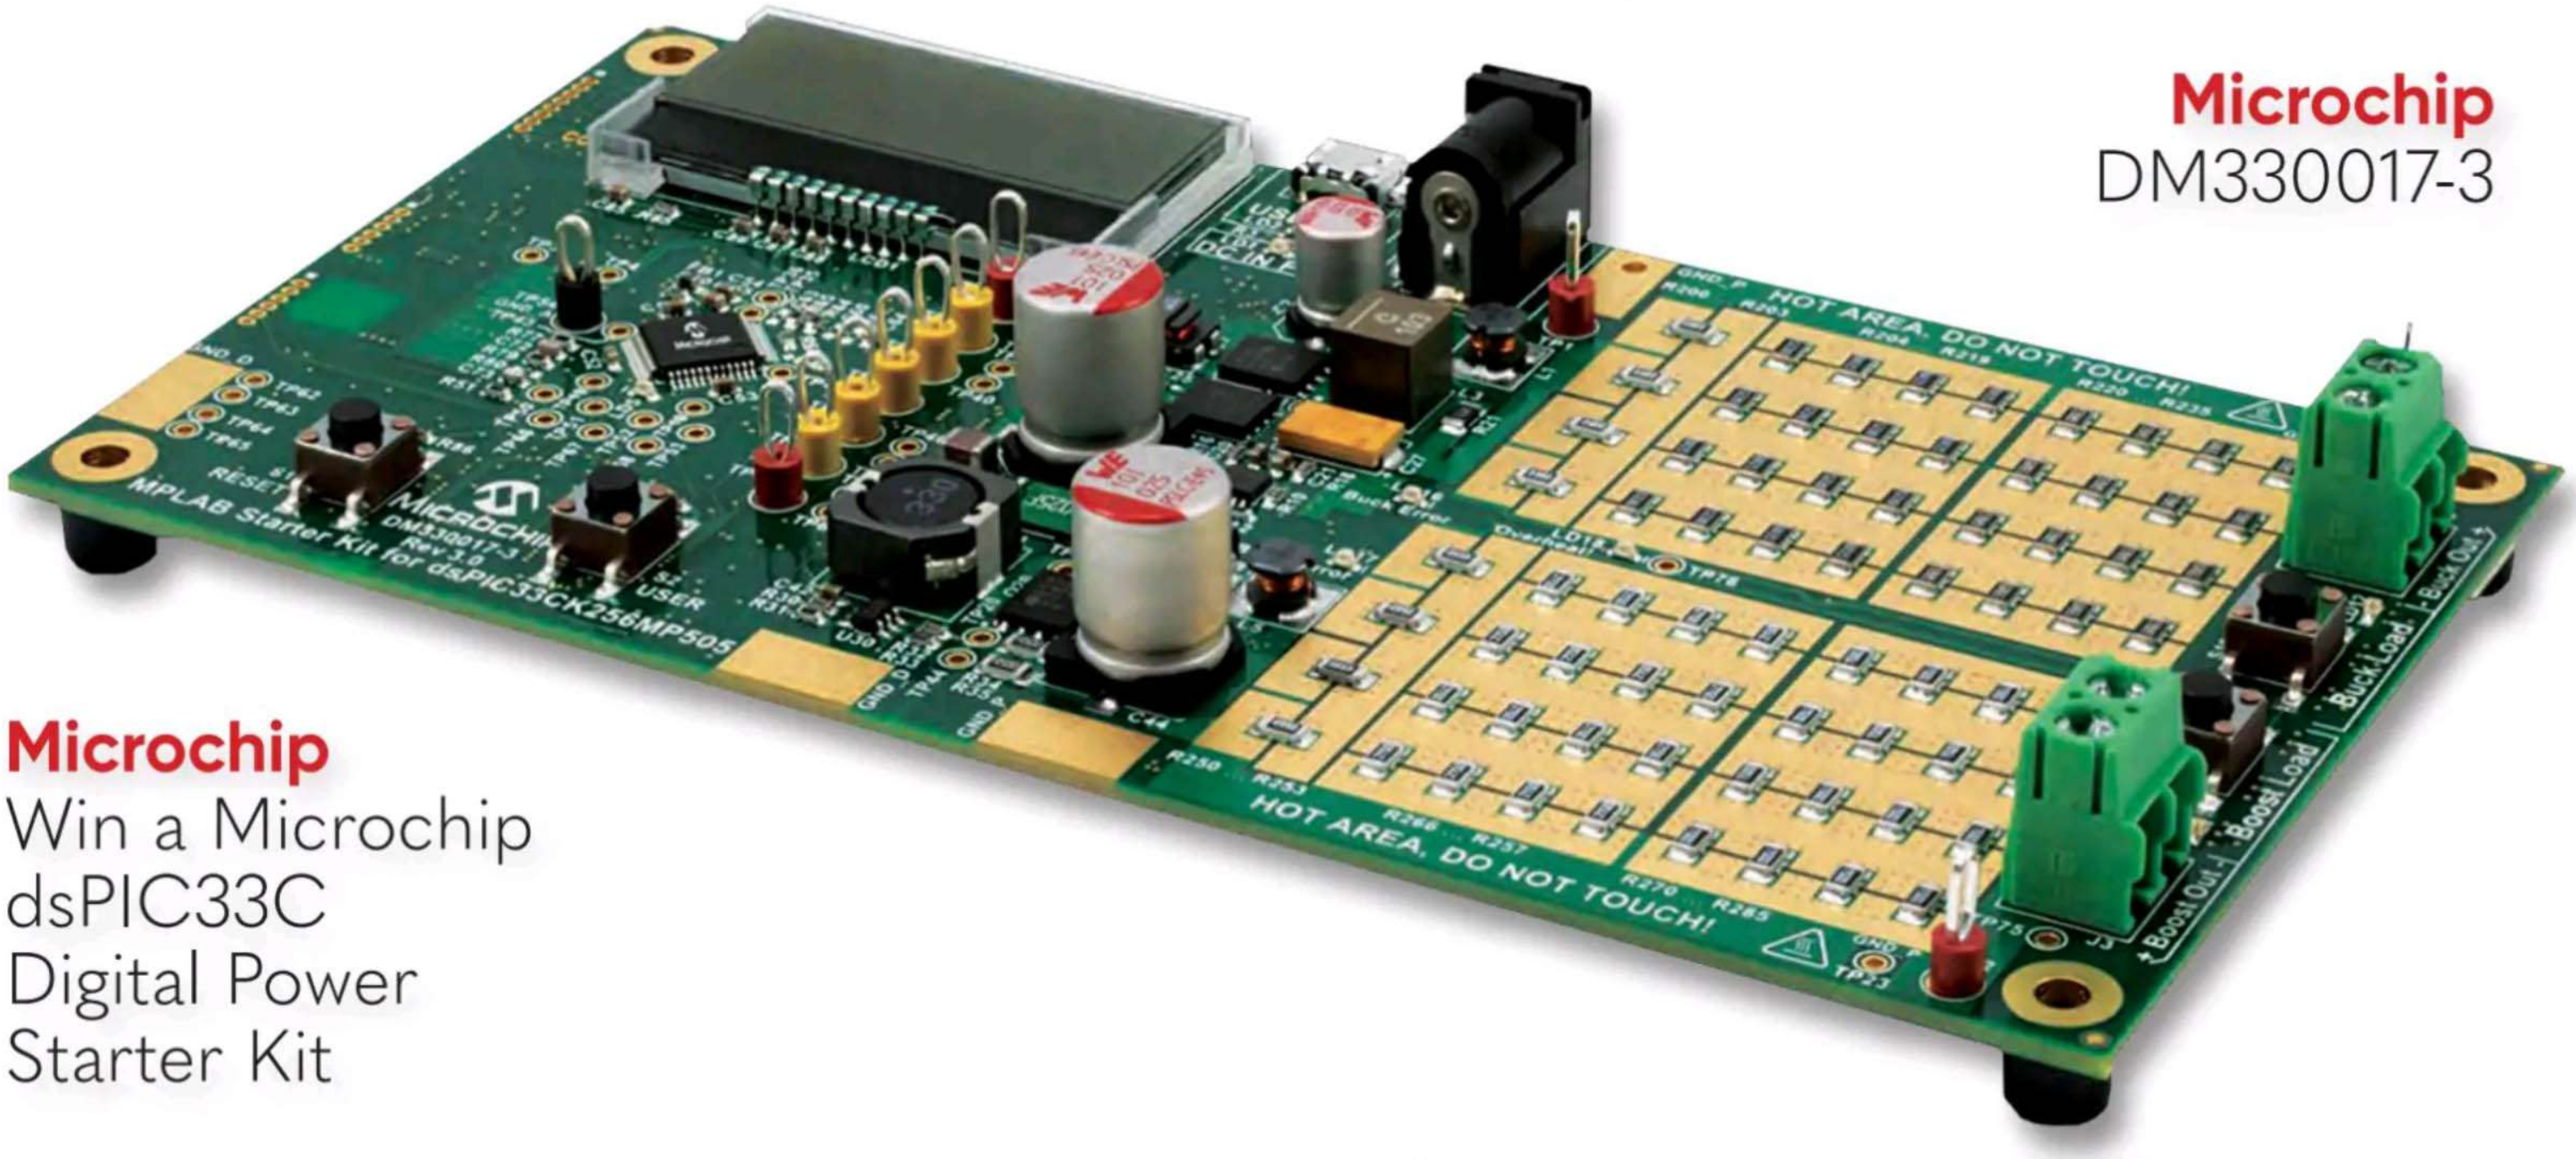

Microchip
DM330017-3

Microchip
Win a Microchip dsPIC33C Digital Power Starter Kit

How to enter

For your chance to win a Microchip dsPIC33C Digital Power Starter Kit or receive a 15%-off voucher, including free shipping, enter your details in the online entry form at:
https://page.microchip.com/PE-33C.html

Closing date

The closing date for this offer is 30 November 2024

October 2024 winner

Jack Booker

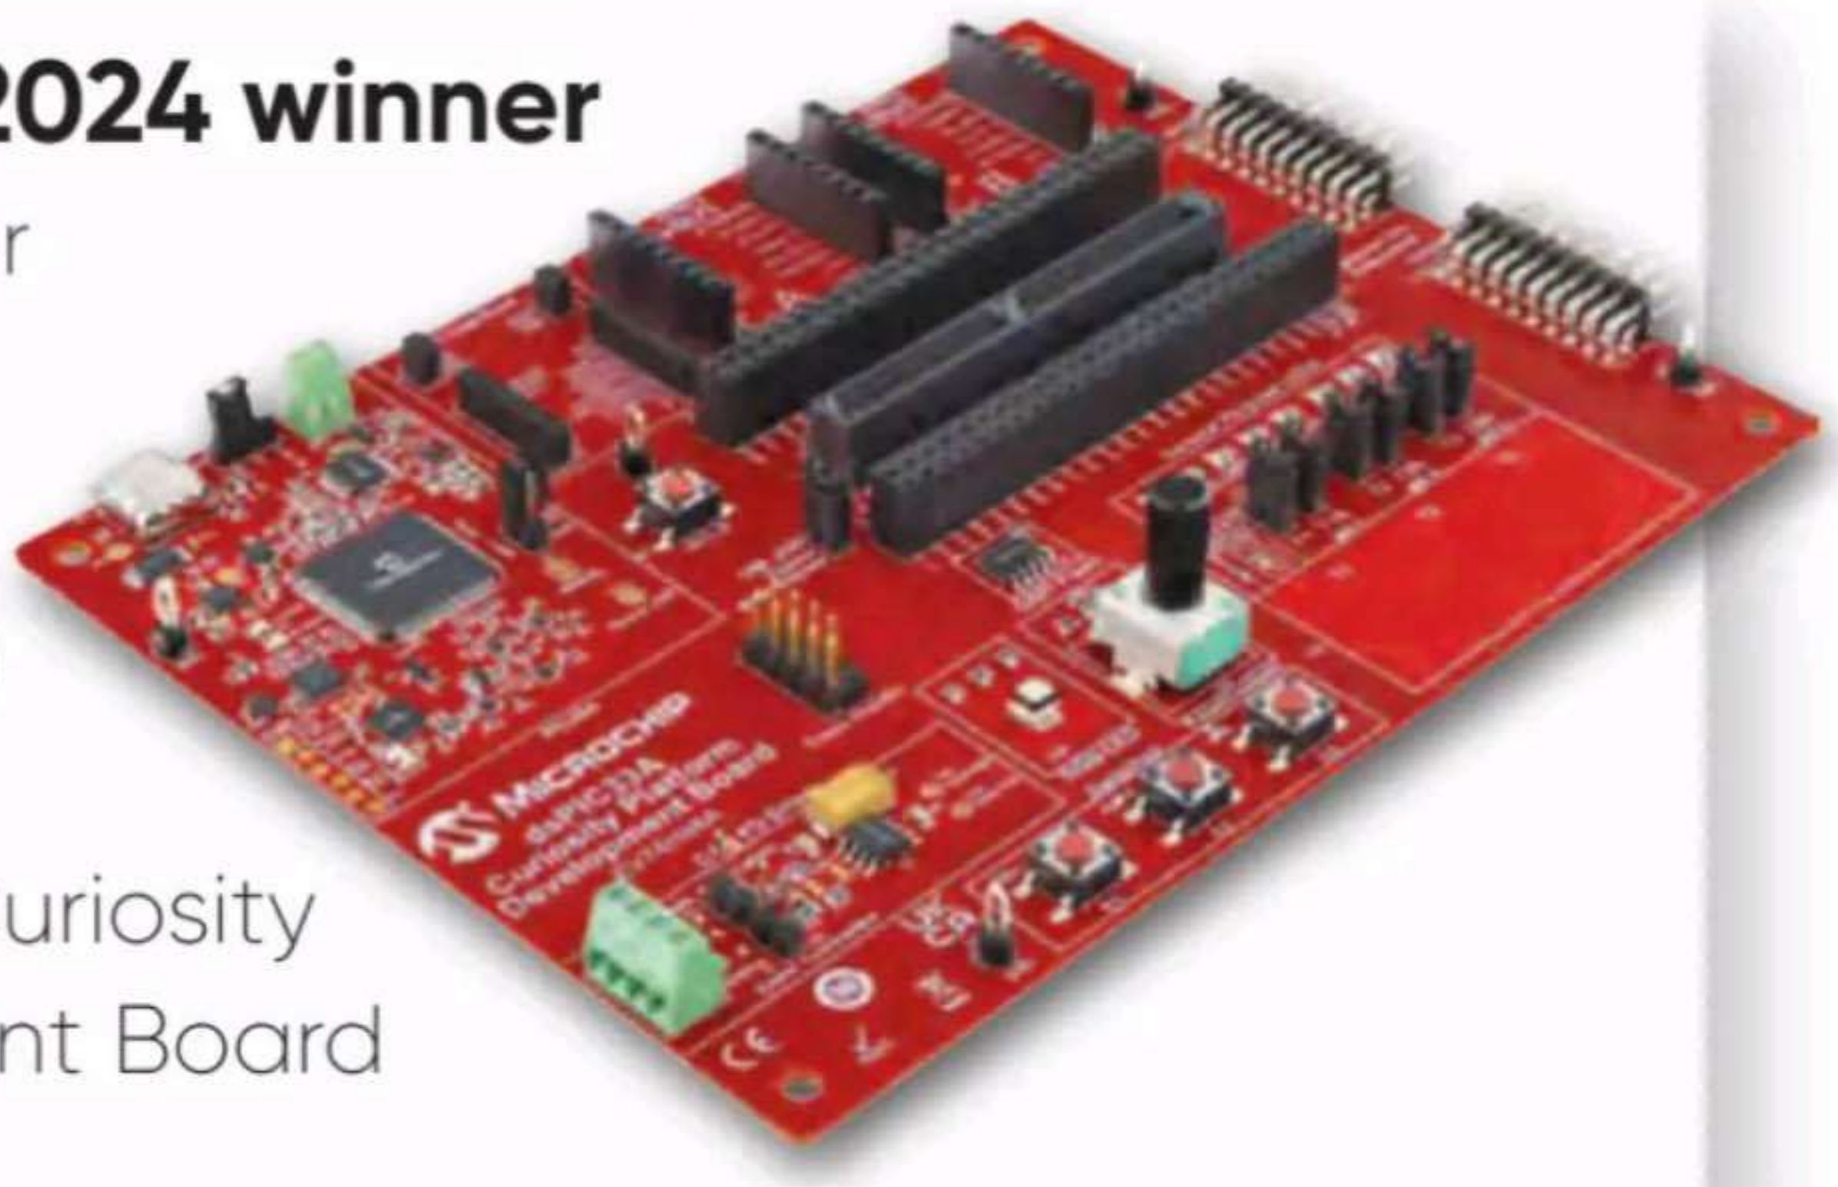

Jack won a Microchip dsPIC33A Curiosity Development Board

VGA PicoMite

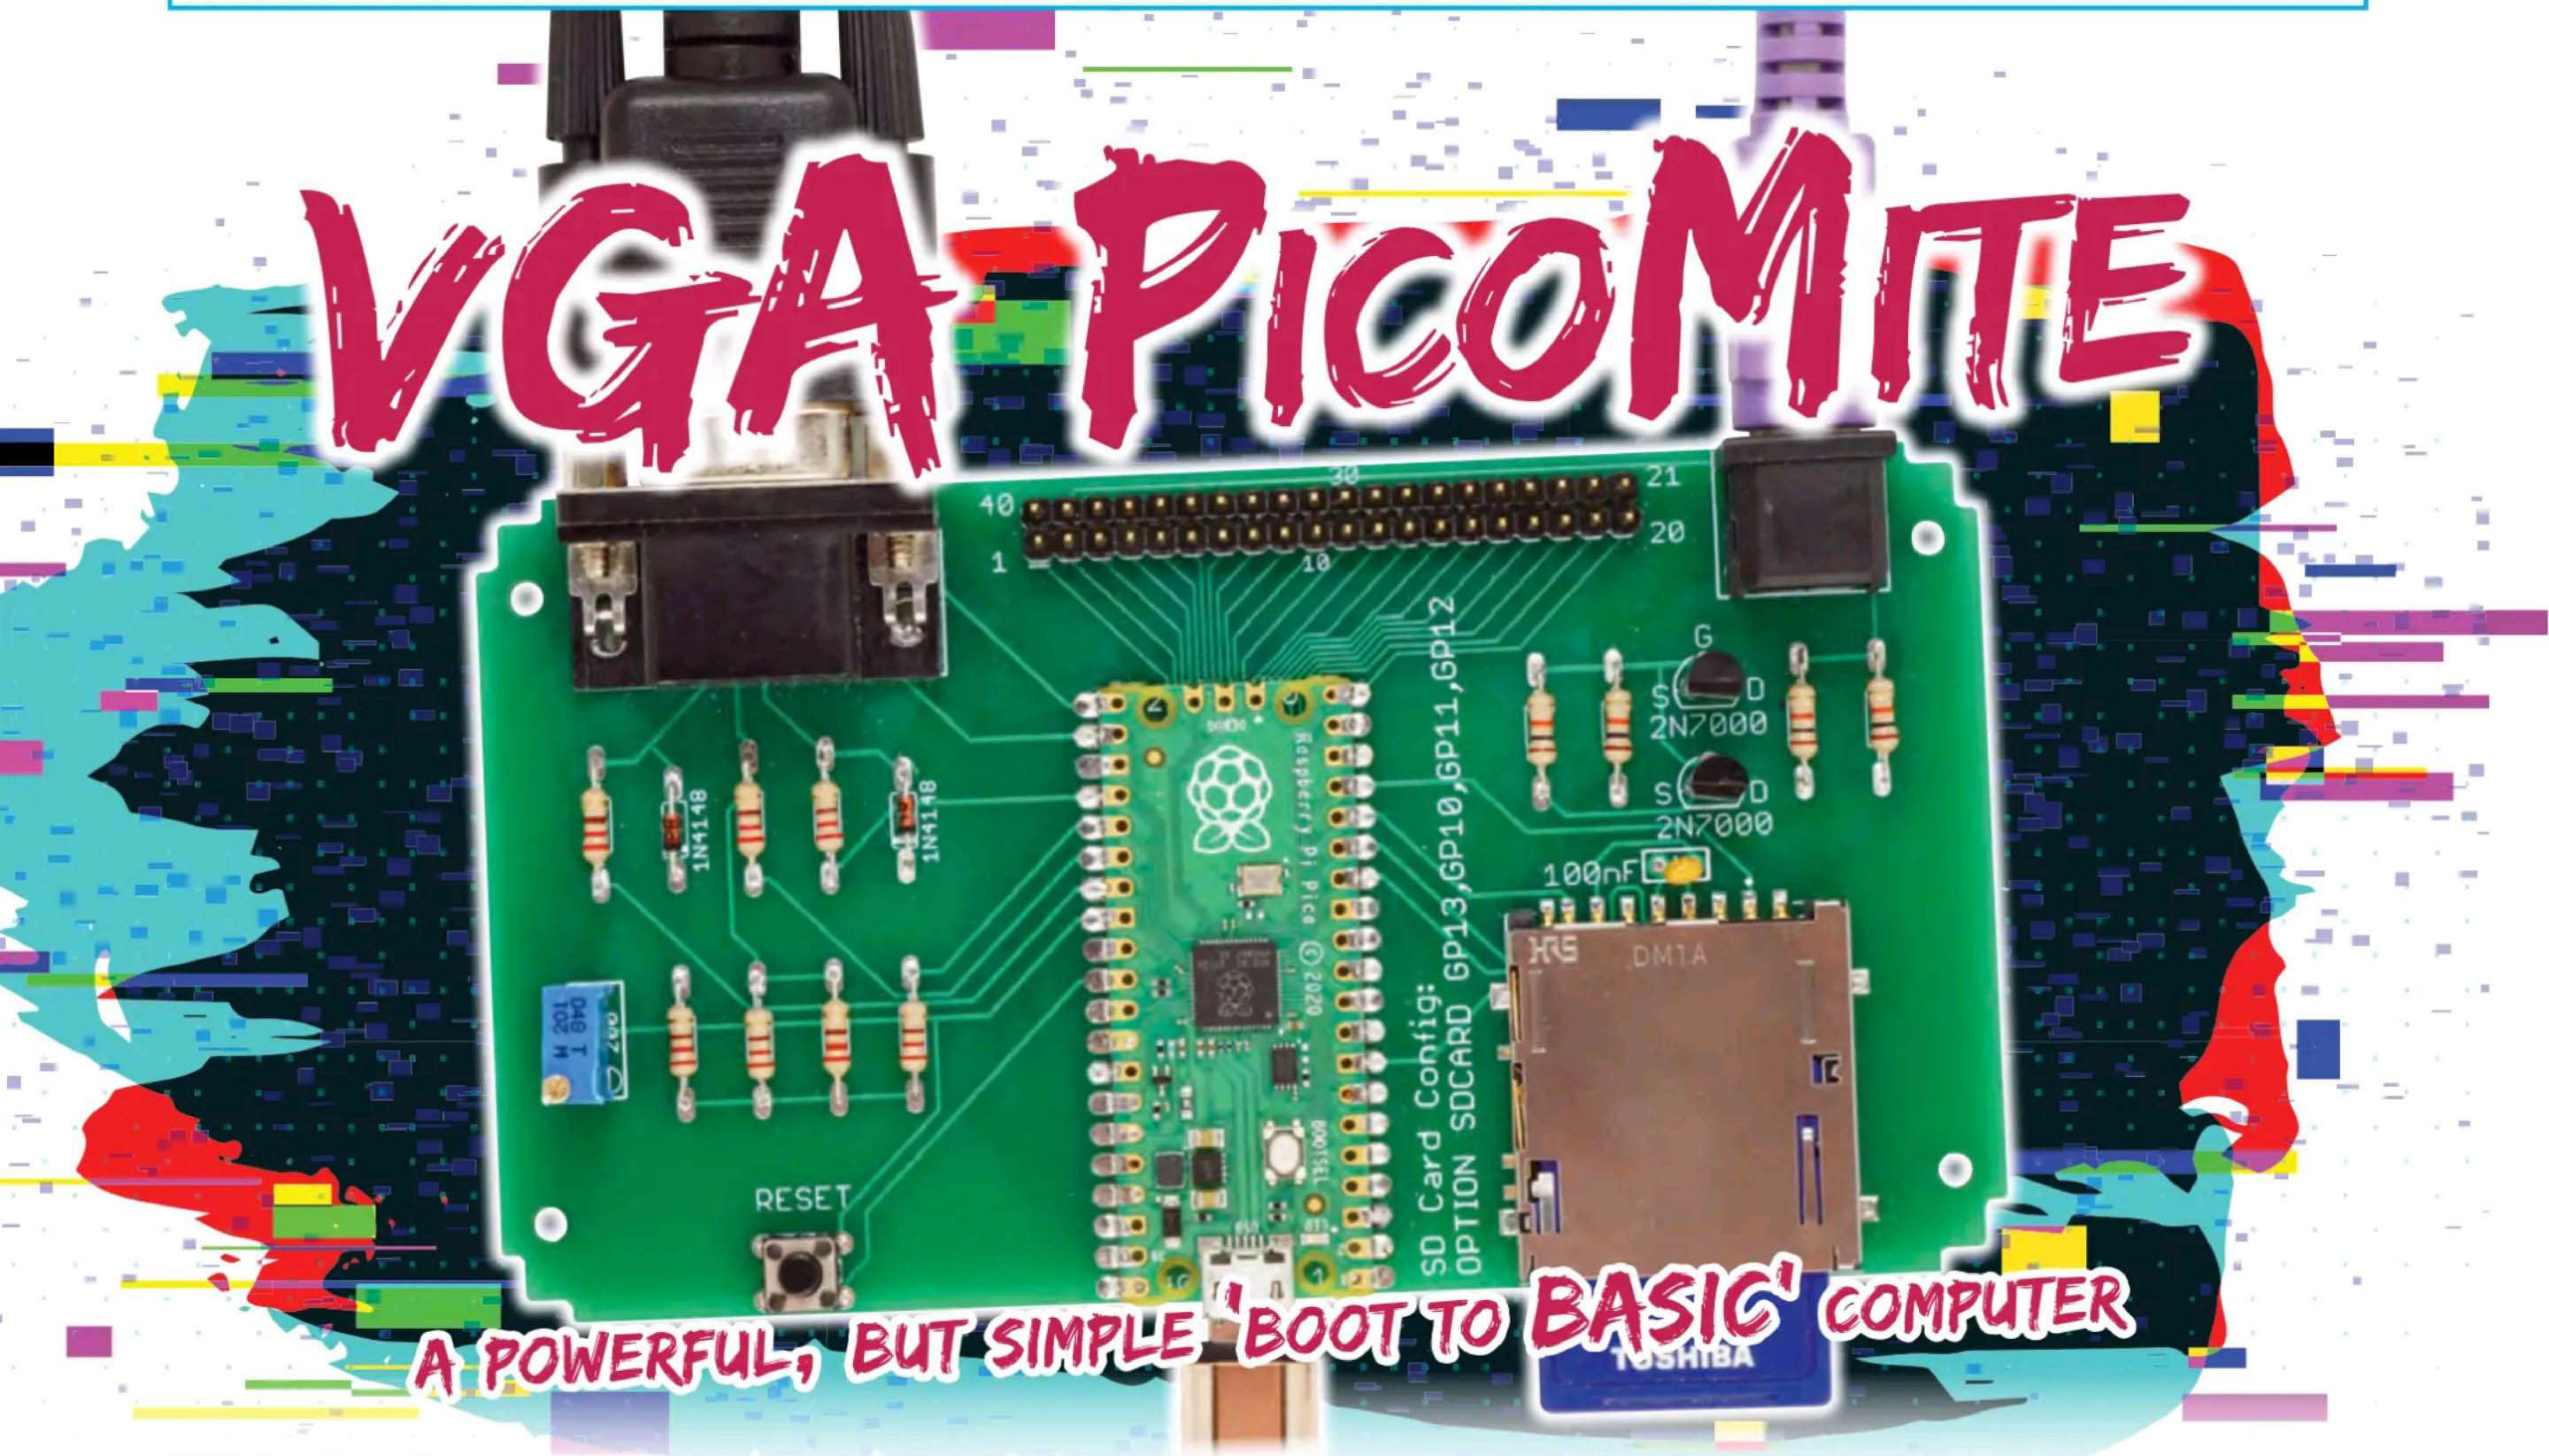

A powerful, but simple 'boot to BASIC' computer

With the low-cost Raspberry Pi Pico board, a handful of standard components and our new firmware, you can build this amazingly capable 'boot to BASIC' computer. It has a 16-colour VGA output, a PS/2 keyboard input, runs programs from an SD card and can be built in no time. And it uses our popular fully-featured MMBasic interpreter.

AFTER THE FUN OF BUILDING IT, you can have even more fun writing programs to track your expenses, teach children about computing, play simple computer games and calculate the positions of the planets (to name a few possible activities).

'Boot to BASIC' computers were popular in the early 1980s – think of the Tandy TRS-80, Vic 20, Commodore 64 and 128, Apple II, Amigas and many others. After turning them on, within a fraction of a second, you have access to the built-in BASIC language, where you can test commands, then enter or load a program and run it. The operating system is built-in.

Developing a program is fast and stress-free. You use the editor to enter the program into memory and run it with one keystroke. If there is an error in your program, you get a plain English message pointing to the problem. With another keystroke, you will be taken back into the editor with the cursor placed on the line that caused the error, ready to be fixed.

This computer is the latest of a series of BASIC-powered computers that have been designed over more than a decade, all referred to as "Maximites", as opposed to Micromites, which are microcontrollers that are also programmed in BASIC.

What sets this one apart is that it is so simple. The Raspberry Pi Pico module costs less than $10, and with just a few other components, you have a fully working computer.

While it is simple, it does not lack performance. It runs a dual-core 32-bit CPU with a clock speed of up to 252MHz, has 103KB of program memory, a 640 × 480 pixel VGA output and stores programs and data on standard SD cards. To put that into perspective, it is about 100 times faster and more capable than the Apple II computer that was so popular in its heyday, yet costs about 1/100th as much!

This firmware is derived from the PicoMite firmware for the Raspberry Pi Pico that we first mentioned in the July 2022 "Make it with Micromite" article on installing MMBasic on a Raspberry Pi Pico.

In this version of the PicoMite, we have removed the ability to drive LCD panels and replaced it with the ability to produce a colour VGA signal. Other than this, the two versions are identical.

VGA output

The VGA output is generated within the Raspberry Pi Pico, a considerable feat for a low-cost processor without any specialised hardware for the task. It is derived from work by Miroslav Nemecek in the Czech Republic, ported by Peter Mather in the UK.

Words and MMBasic by Geoff Graham   VGA PicoMite firmware by Peter Mather

VGA PicoMite

The Raspberry Pi Pico (shown 50% larger than real life) is a popular and cheap microcontroller module with plenty of memory, speed and I/O capability. With our VGA PicoMite firmware, you can easily program it in BASIC with a VGA monitor, keyboard and SD card storage support and access to all the Pico's features.

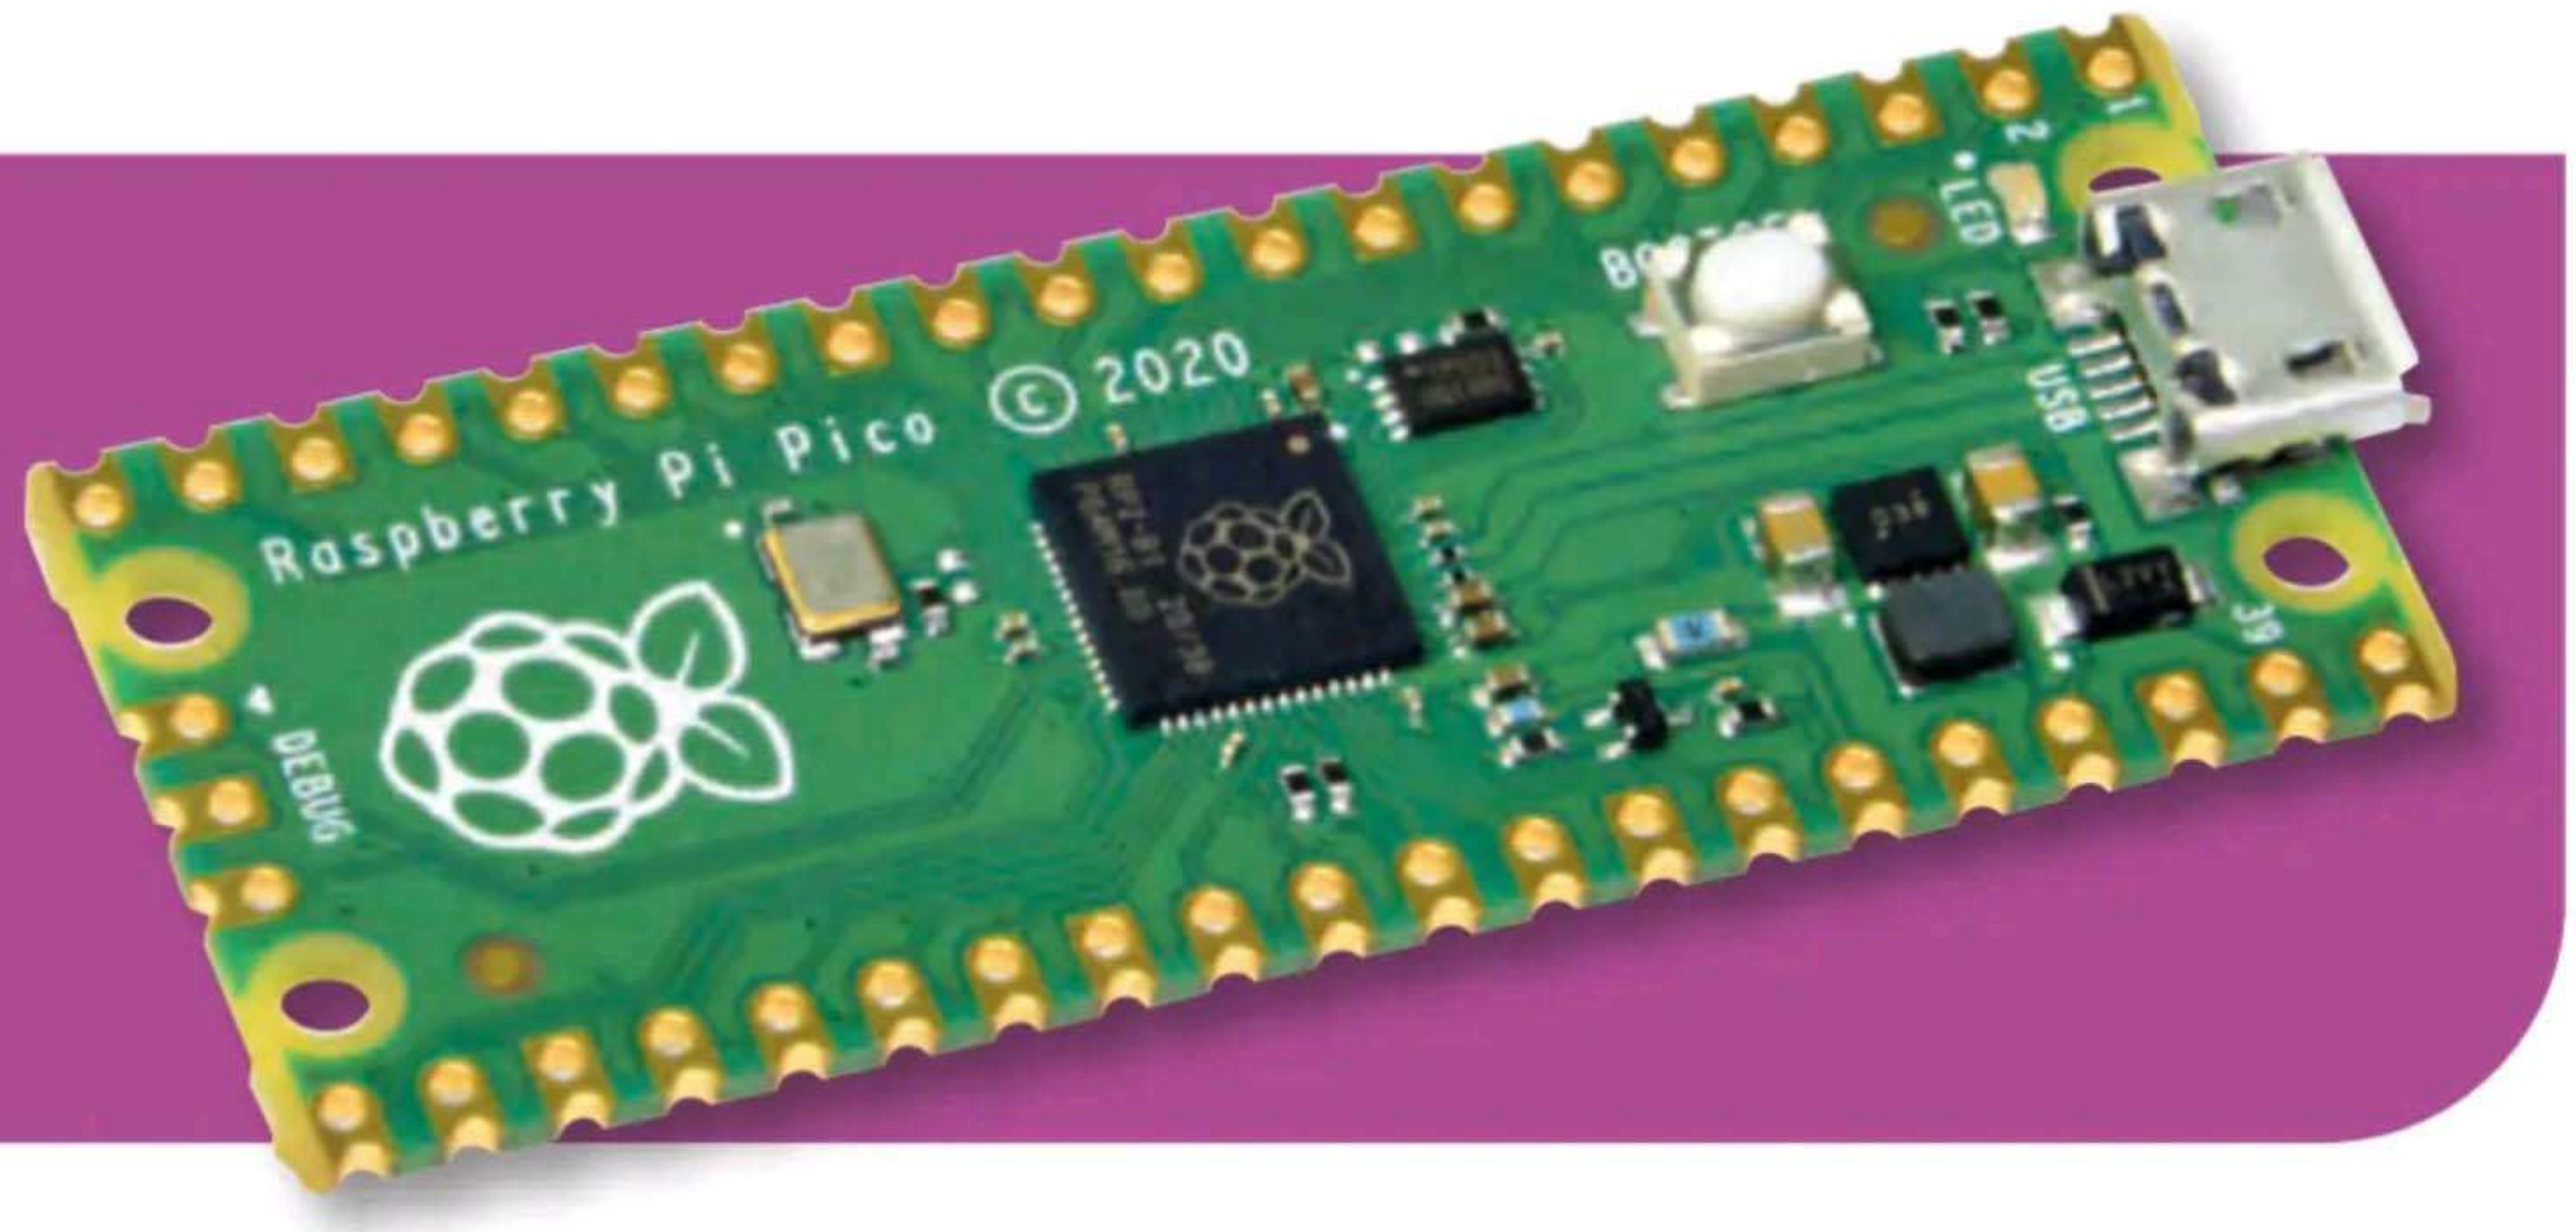

The VGA signal is generated using the second CPU in the RP2040 processor on the Raspberry Pi Pico plus one PIO channel. Because of this, it does not affect the performance of the BASIC interpreter, which runs unimpeded at full speed on the first CPU.

The video signal is in the standard VGA format with a pixel rate of 25.175MHz and a frame rate of 60Hz, so you get a stable video output that is supported by all VGA monitors. In monochrome mode, the resolution is 640 x 480 pixels, while in colour mode, the pixels are doubled horizontally and vertically to give an effective resolution of 320 x 240 pixels.

Regardless of the mode, the monitor still receives a 640 x 480 pixel image, and virtually all monitors support this mode.

The colour mode uses four bits to define the colour of each pixel, giving a total of sixteen (2^4) colours (including black). One bit is assigned to red, two bits for green and one bit for blue. While this is not very good for displaying photographs, it is great for games, graphs and brightening up the output of your programs.

The graphics driver includes some of the features of the Colour Maximite 2. So, in addition to drawing lines, boxes, etc, you can also have multiple fonts and create fast animations with the BLIT command. You can load images from the SD card for display on the video screen, and you can save an image of the current screen to the SD card (ie, take a screenshot/screengrab).

The hardware to generate the VGA signal only uses eight resistors and two diodes to generate the correct signal levels. These components can be assembled on a piece of perforated stripboard (eg, Veroboard), but we felt that many people would like the convenience of building it on a PCB, so we have designed a simple board that includes the VGA output and the keyboard and SD card interfaces.

But, if you want to build it on an extreme budget, the stripboard solution will work fine.

Keyboard and SD card

The VGA PicoMite firmware supports standard PS/2 keyboards. While they are less popular today than USB keyboards, there are still plenty to be found, and many wired USB keyboards come with a PS/2 adaptor. These work perfectly with the VGA PicoMite. The keyboard is fully supported; auto-repeat and the arrow and function keys work as expected.

The PS/2 standard uses 5V signal levels, but the Raspberry Pi Pico is strictly limited to a maximum of 3.6V, so level shifting is required to interface the Pico with the keyboard. This is done using four resistors and two Mosfets. Like the VGA interface, this could be built on a perforated stripboard if you do not want to use our PCB (although it isn't exactly expensive).

SD cards use 3.3V signal levels, so they can connect directly to the Raspberry Pi Pico with no interface components required. The firmware supports cards up to 32GB, formatted with a FAT16 or FAT32 file system. MMBasic supports long file names, hierarchical folders etc. You can save, load and run BASIC programs from the card as well as read/write data files to it.

The file system is compatible with Windows, Linux and macOS, so it is easy to use an SD card to transfer data back and forth between the VGA PicoMite and a computer or laptop.

However, an SD card is not necessary as the firmware reserves eight 'slots' for BASIC programs in the Pico's flash memory. So you can save or retrieve up to eight separate programs there if you wish. If you are just casually playing with this computer, that could be all you need.

Regardless, adding SD card support is so easy that we have included it on our PCB.

The firmware supports many other devices that can be connected, including infrared (IR) remote control receivers, stereo audio output, distance sensors, a real-time clock, temperature sensors and so on. We did not clutter up the PCB with these, but you can easily connect external circuitry to the 40-pin header on the rear of the board.

These pins mimic the I/O pins on the Raspberry Pi Pico (including the VGA and keyboard connections), so you can add almost whatever you want via that connector.

MMBasic

None of this can work without the firmware and MMBasic. We have described MMBasic many times before, so it is sufficient to say that it is a full-featured version of the BASIC computer language designed to be easy to use.

As mentioned before, the firmware includes a built-in editor with a colour-coded display plus support for all the peripherals described above (and more).

On the VGA PicoMite, programs can be as large as 108KB. The amount of RAM available for buffers, arrays, etc is even larger at 140KB, so you can have very large data arrays.

By default, the CPU runs with an instruction clock of 126MHz, but you can switch it to 252MHz. This is 'overclocking' the RP2040 processor used in the Raspberry Pi Pico, which has a specified top speed of 133MHz. However, nearly all the Picos we have tested run fine at this speed, so it is a viable option if you want to go mad with the performance.

One feature that will be appreciated by people who played with the early computers of the 1980s is that MMBasic saves the program to flash memory, not RAM. When you edit a program, it is saved to flash, and when you run it, it runs from flash.

Because flash memory is non-volatile,

you will not lose your program, even if you accidentally restart the device or interrupt the power – something that happened distressingly often with the early computers that stored programs in volatile RAM.

Circuit details

Fig.1 is the straightforward circuit of the VGA PicoMite. The VGA output from the Raspberry Pi Pico uses six digital signals – one for red, two for green, one for blue and one each for horizontal and vertical synchronisation. For the red and blue signals, the digital output levels are simply clipped by the 1N4148 diodes to 0.7V, the correct level for full brightness in the VGA standard.

The green output also has a resistor network providing four intensity levels from completely off to full brightness. The 200Ω multi-turn trimming potentiometer lets you adjust for a balanced white output without needing precision resistors with hard-to-find values.

The horizontal and vertical sync signals use standard TTL signal levels, so they directly connect to the monitor.

The keyboard interface consists of just two signals, the clock and the data lines, which are bi-directional. Each signal is level-shifted using two resistors and a small 2N7000 Mosfet. There are many ways to implement level shifting, but this is cheap and uses common components.

In the idle state, the outputs from the Pico and the keyboard are pulled high to 3.3V and 5V respectively, by 10kΩ pull-up resistors. When the Pico wants to send a signal, it pulls its output low, causing the Mosfet to conduct because its gate is held at +3.3V. This means that the keyboard's corresponding pin is also pulled low.

When the keyboard wants to signal, it pulls its pin low, and that causes the substrate diode in the Mosfet to conduct and also pull the Pico's corresponding pin low.

There is very little to the SD card interface. All the SD card signal lines connect directly to the Raspberry Pi Pico as both the SD card and the Pico use 3.3V logic signalling.

This SPI serial bus consists of four signals:

- Chip select ($\overline{CS}$), which is pulled low by the Pico when it wants to communicate.
- The clock, generated by the Pico.
- Master out slave in (MOSI), which carries data from the Pico (the master) to the card.
- Master in slave out (MISO), which carries data from the card to the Pico.

The card socket also has switches to indicate when a card is inserted and whether it is write-protected, but these are not used to keep the maximum number of the Pico's I/O pins free. Instead, the card's insertion and removal are detected by 'polling' the SD card (ie, periodically checking if it is present).

The Raspberry Pi Pico itself is powered via its micro USB connector, which provides 5V for the PS/2 keyboard.

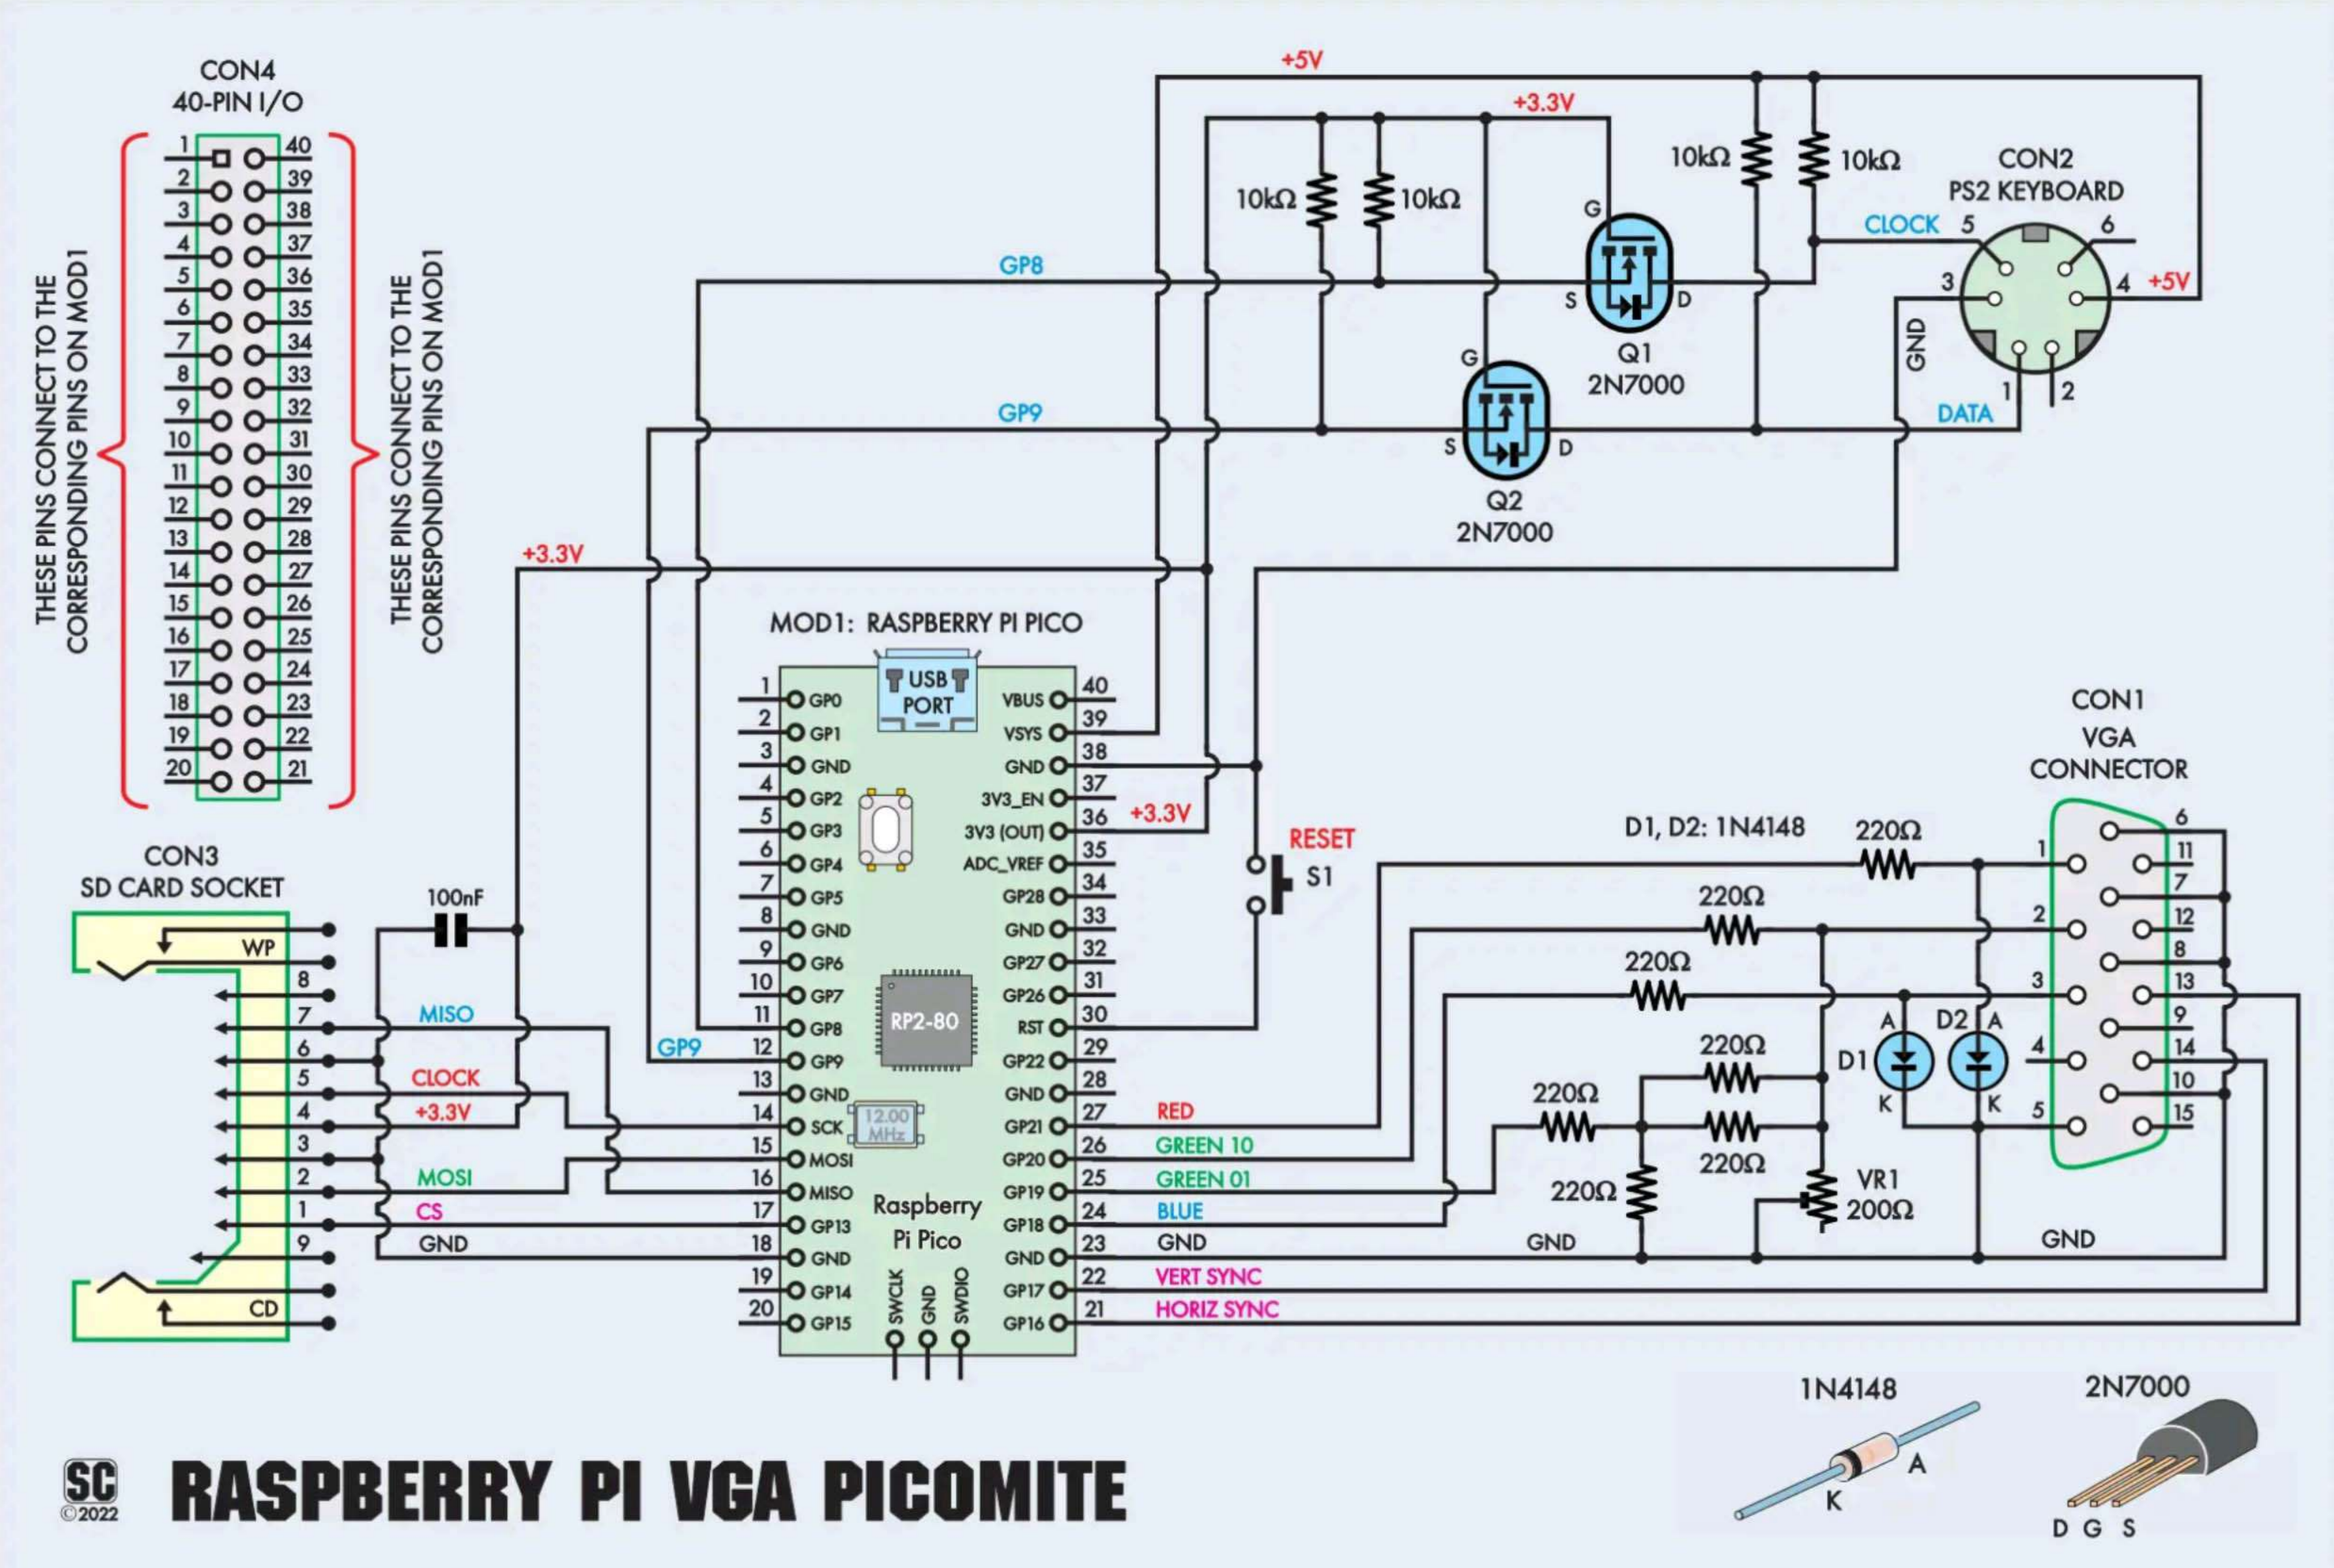
CON4
40-PIN I/O
THESE PINS CONNECT TO THE CORRESPONDING PINS ON MOD1
+5V
+3.3V
10kΩ
10kΩ
10kΩ
10kΩ
CON2
PS2 KEYBOARD
CLOCK
DATA
GND
GP8
GP9
Q1
2N7000
Q2
2N7000
G
S
D
MOD1: RASPBERRY PI PICO
USB PORT
RP2-80
12.00 MHz
Raspberry Pi Pico
GP0
GP1
GND
GP2
GP3
GP4
GP5
GND
GP6
GP7
GP8
GP9
GND
SCK
MOSI
MISO
GP13
GND
GP14
GP15
SWCLK
SWDIO
VBUS
VSYS
3V3_EN
3V3 (OUT)
ADC_VREF
GP28
GP27
GP26
RST
GP22
GP21
GP20
GP19
GP18
GP17
GP16
RESET
S1
CON3
SD CARD SOCKET
WP
CD
100nF
MISO
CLOCK
MOSI
CS
RED
GREEN 10
GREEN 01
BLUE
VERT SYNC
HORIZ SYNC
D1, D2: 1N4148
220Ω
D1
D2
A
K
VR1
200Ω
CON1
VGA CONNECTOR
1N4148
2N7000
D G S
SC ©2022
RASPBERRY PI VGA PICOMITE

Fig.1: the VGA PicoMite circuit mainly comprises the components necessary to convert the Raspberry Pi Pico's signal levels to that required by the PS/2 keyboard and VGA monitor. It could be assembled on a piece of perforated stripboard, but the commercial PCB doesn't cost that much, makes construction much easier and gives a more professional-looking result.

The output of the Pico's onboard 3.3V regulator is used to power the processor and flash memory, and is also fed to the SD card.

For maximum flexibility, all 40 pins on the Raspberry Pi Pico are routed to the 40-way connector on the rear of the PCB in the same configuration as that used by the Pico. This makes it easy to connect external devices using jumper wires, as you can consult the Pico pinout diagram and then select the corresponding pins on the 40-way connector.

Taking into consideration the I/O pins reserved for the VGA output, keyboard and SD card, there are still 14 I/Os available for interacting with external circuitry.

Construction

The simplest way to source the components is to purchase a kit from the Silicon Chip Online Shop, which includes everything except a power supply and a case. But if you decide to source the parts yourself, you might have some difficulties due to current shortages, especially with the VGA and SD card connectors.

It can help to use a parts search site like https://octopart.com to identify suppliers that currently have stock of a given part.

Populating the PCB is straightforward (see Fig.2). As usual, start with the low-profile components and work upward. Preferably, your soldering iron should be temperature-controlled and have a chisel or conical tip with a diameter of 1.7mm or thereabouts.

You can get away with other sizes, but some of the pads (for example, on the VGA connector) are close together, meaning a smaller tip will be easier to use.

Note the choice of mounting a vertical or right-angle 2x20-pin header for CON4 depending on whether you're using a case or not.

The SD card socket is surface-mounted and, because you need space to get your soldering iron close to it, you should start with that. Begin by applying a thin layer of flux paste on all its pads. It has two small posts on the underside that click into matching holes in the PCB to ensure perfect alignment.

With the socket in position, solder the two tabs on the right side of the socket (viewed from the front) and the five on the left side. Now that the

Parts List – VGA PicoMite

1 double-sided PCB coded 07107221, 124mm x 69mm
1 Raspberry Pi Pico
1 right-angle PCB-mount DE15 (VGA) socket (CON1) **[TE Connectivity 1-1734530-1 or Multicomp SPC15430; Mouser, Farnell etc]**
1 right-angle PCB-mount 6-pin mini-DIN socket (CON2) **[Farnell 1200113]**
1 Hirose DM1AA-SF-PEJ(72) or DM1AA-SF-PEJ(82) SD card connector (CON3) **[Mouser, DigiKey, Farnell etc]**
1 2×20-pin header, 2.54mm pitch (CON4) ***OR***
1 2×20-pin right-angle box header, 2.54mm pitch (CON4) (for installation in a case)
1 4-pin vertical short actuator tactile switch (S1) **[Farnell 4511189]**
1 200Ω multi-turn top-adjust trimpot (VR1) **[Farnell 9353569]**
2 2N7000 60V 200mA N-channel Mosfets, TO-92 (Q1, Q2)
2 IN4148 diodes (D1, D2)
1 100nF 50V+ multi-layer ceramic or MKT capacitor

Resistors *(all through-hole ¼W 5%)*
4 10kΩ metal or carbon film resistors
7 220Ω metal or carbon film resistors

Optional parts
1 71 × 130 × 30mm grey ABS instrument case **[Altronics H0376]**
4 No.2 × 4-5mm self-tapping screws (if mounting PCB in case)
4 stick-on rubber feet or short tapped spacers with M3 machine screws (if using PCB without case)
2 20-pin SIL headers, 2.54mm pitch (to make Pico module pluggable)
2 20-pin SIL header sockets, 2.54mm pitch (to make Pico module pluggable)

Kit: a mostly complete kit for the VGA PicoMite is available from the Silicon Chip Online Shop (Cat **SC6417**) for $37.50+VAT (about £23) + P&P. It includes the PCB and everything that mounts on it. You just need to add a USB power supply, keyboard, monitor and optionally an SD card. This can be mounted in a case although you'll also need a different header for CON4 and some self-tapping screws. There is a version of the kit that comes with the case, Cat **SC7402**, for AU$55.00+VAT (about £34) + P&P.

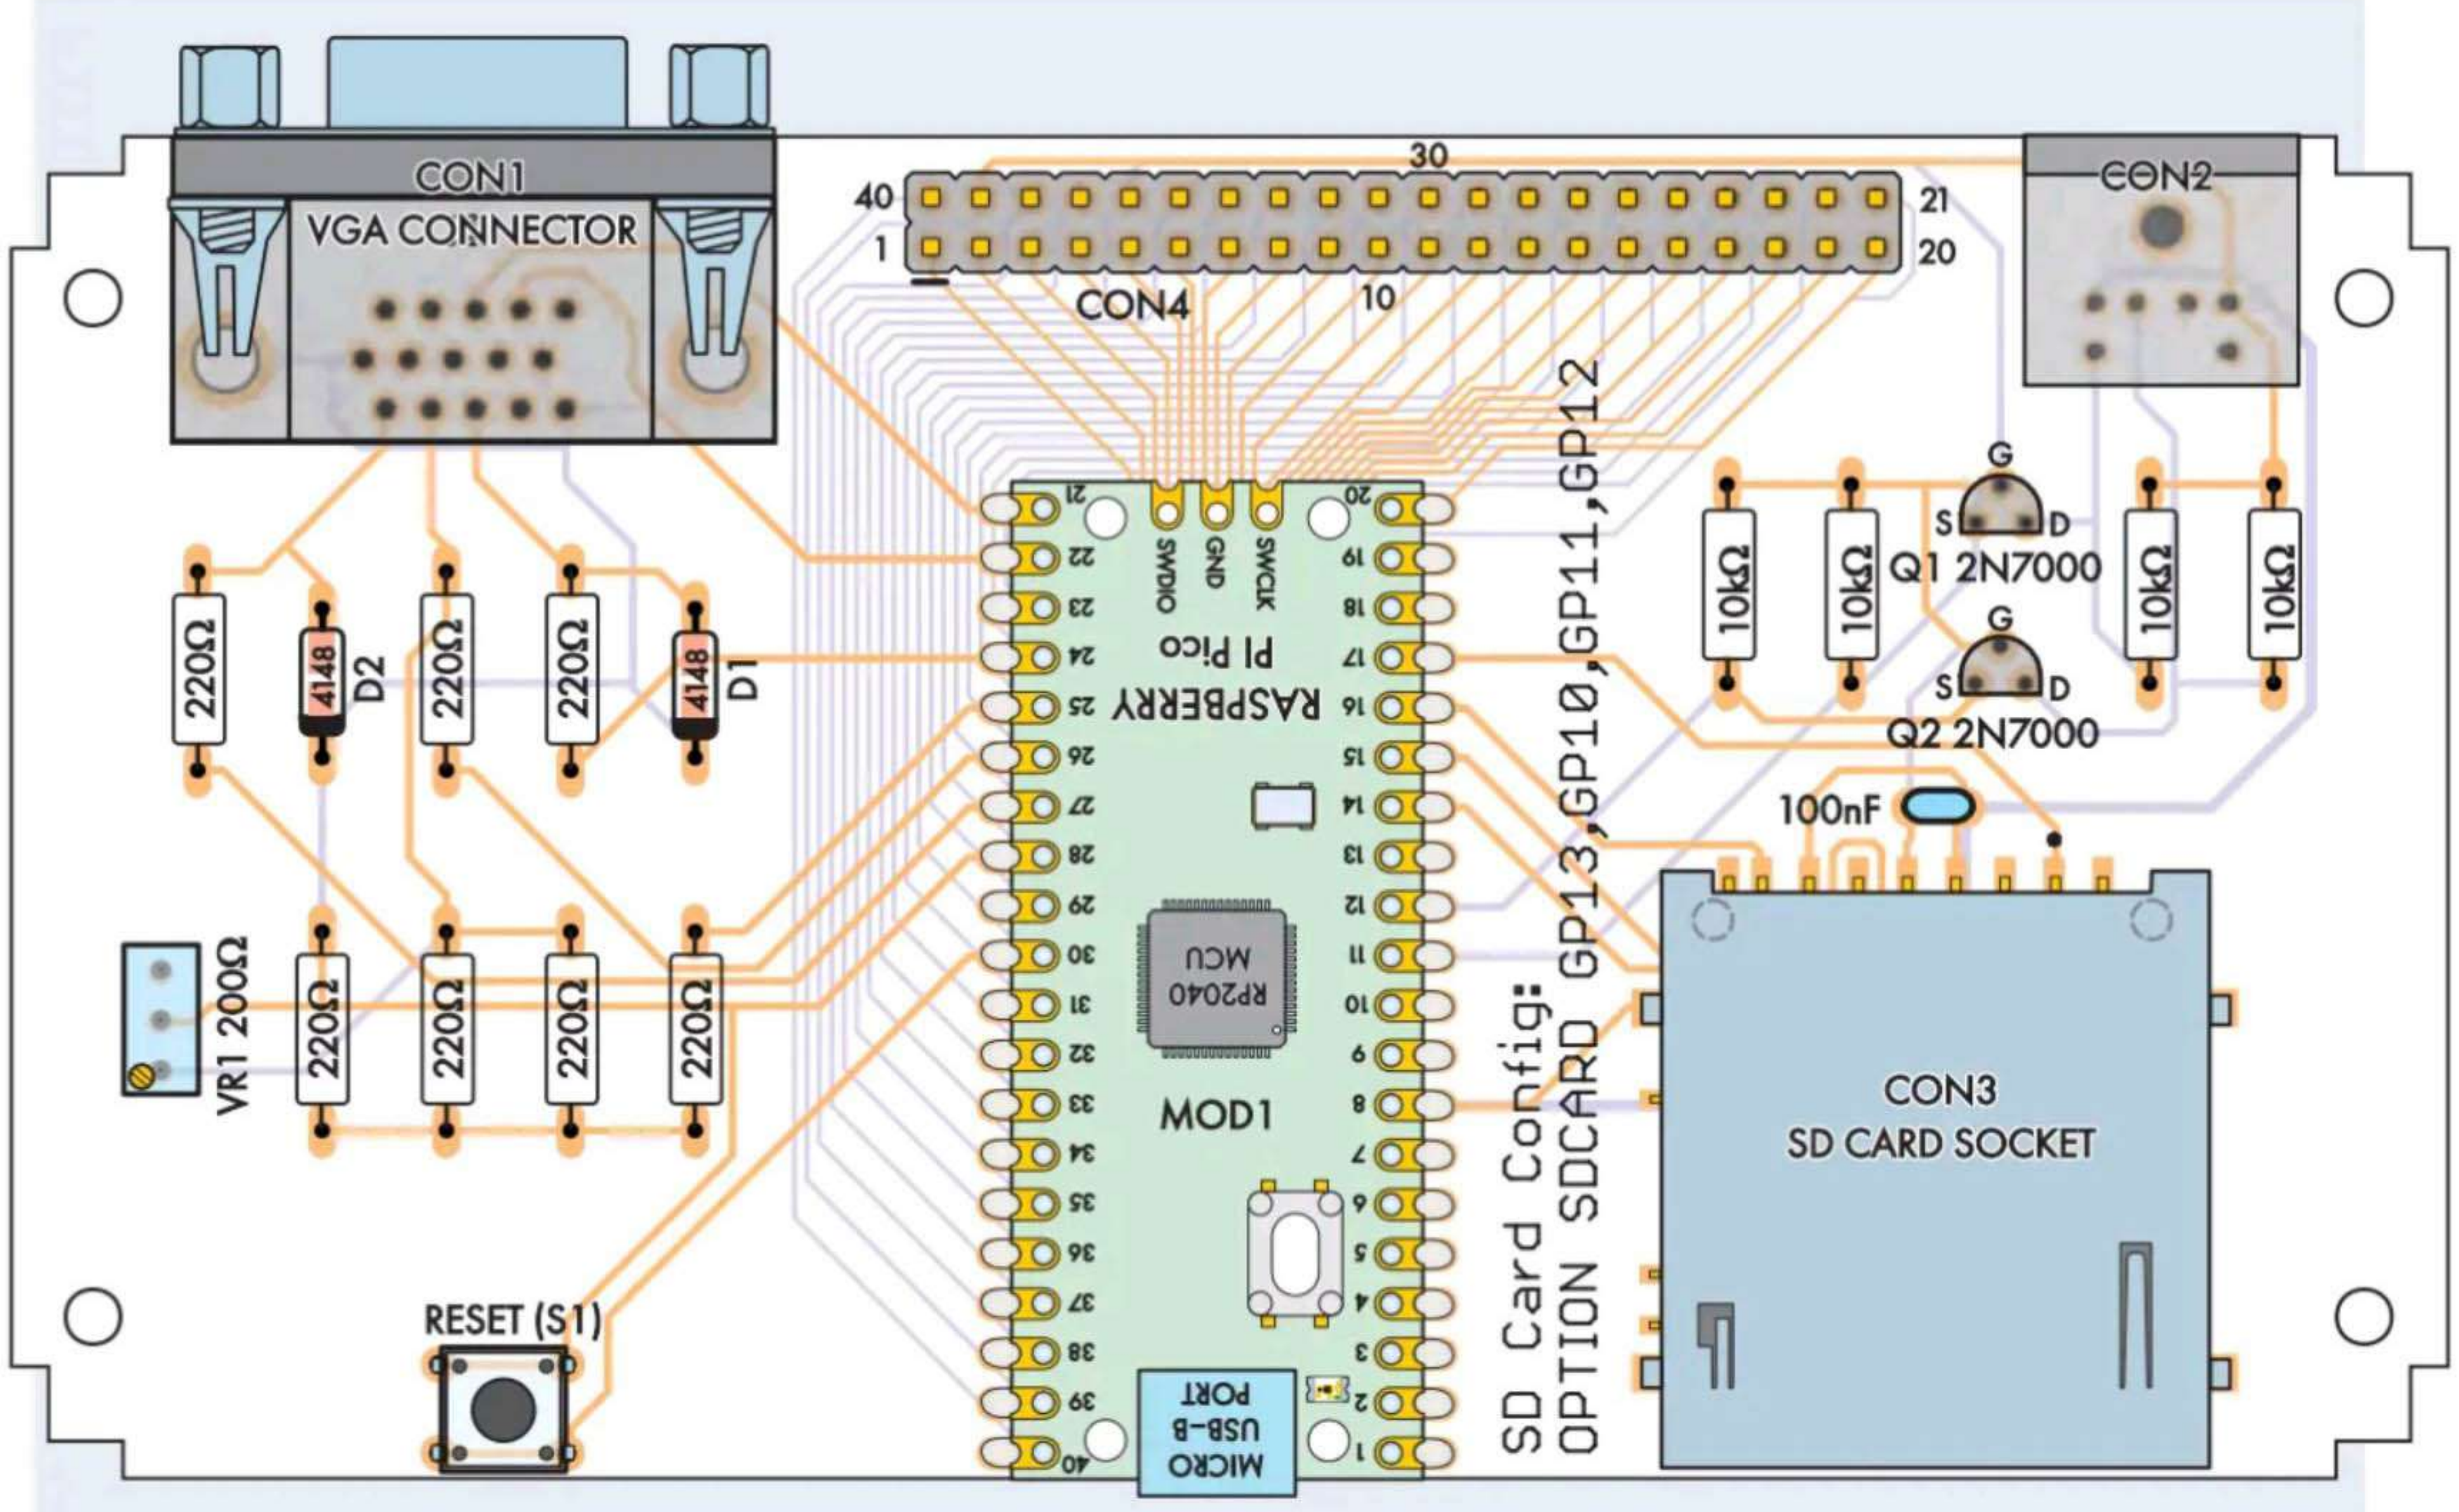

Fig.2: assembling the VGA PicoMite is easy; just fit the parts as shown here, starting with the lower-profile components and working up to the taller ones. The diodes, Mosfets & Pico must be orientated correctly. The main options to consider are whether you're using headers for the Pico and the type of connector you're using for CON4.

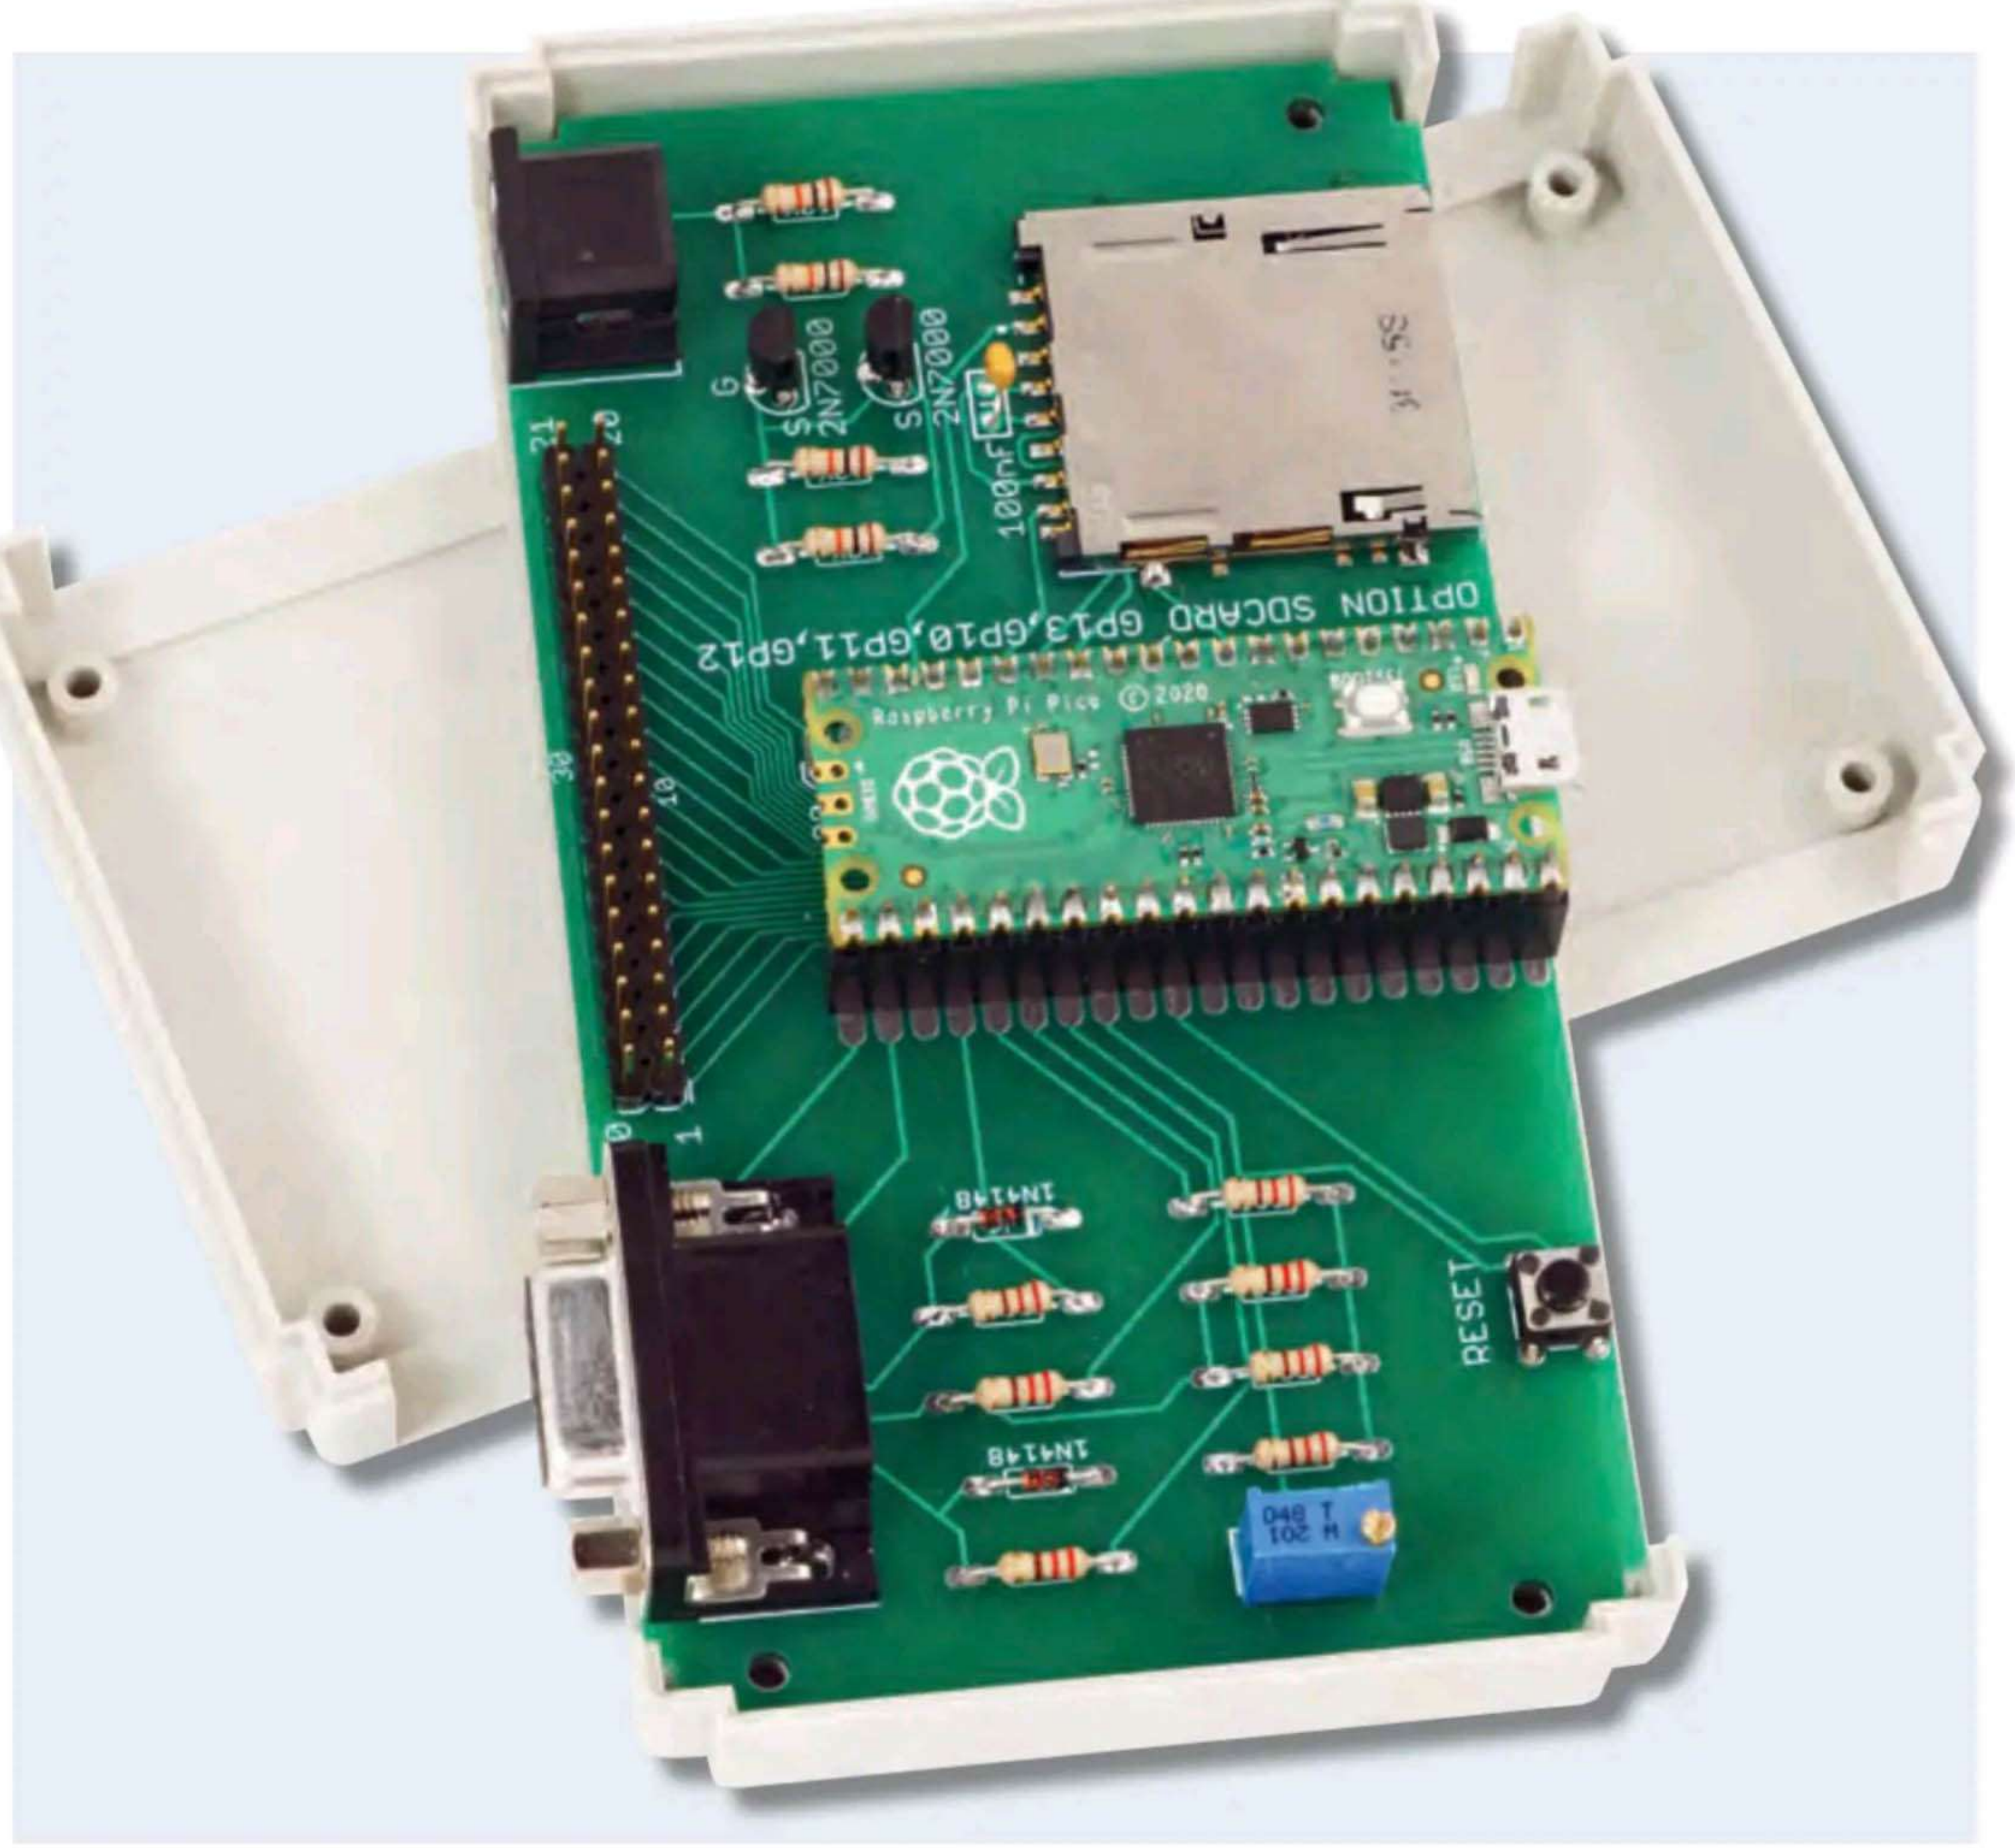

If you want to house the VGA PicoMite in a box, the board is sized to fit in the Altronics H0376 snap-together case. The Raspberry Pi Pico can be directly soldered to the PCB, or you can solder pin headers to the Pico and then plug this assembly into two matching 20-pin single-in-line (SIL) header sockets installed on the PCB. Using headers means that you can easily replace the Pico if it fails.

socket is secured, you can solder the nine pins on the rear. For each pin, slide the tip of your iron over the solder pad towards the connector so that the tip hits the connector's pin. Within half a second, the solder should magically flow around the pin, and you can withdraw the iron.

If you get a solder bridge, don't worry and carry on with the other pins. Finally, examine your soldering using a powerful magnifier and check for any bridges (especially to the connector's shield) and remove them using more flux paste and solder wick.

Be careful here, as the solder wick can suck up all the solder, so you should recheck the joint after using it (although it usually leaves a sufficient amount behind).

The Raspberry Pi Pico can be directly soldered to the PCB, or you can solder pin headers to the Pico and then plug this assembly into matching SIL header sockets soldered to the PCB, as shown in the adjacent photo. Using the sockets is the safest option as you can then easily replace the Pico if you suspect that it has been damaged.

Housing the VGA PicoMite

We expect that most readers will put rubber feet on the bottom of the PCB and use it as is ('naked'). Another option for feet is tapped spacers via the mounting holes.

Still, if you want to house it in a box, the board is sized to fit in a small 71 x 130 x 30mm instrument case (such as Altronics' H0376, made by Ritec). The PCB design has cutouts in the corner that are designed to clear the steps in the corners of that Ritec case.

If you are planning to house the PCB in this, you can replace the vertical 40-way pin header on the rear of the PCB with a right angle shrouded IDC connector so that you can use a ribbon cable to connect to any external circuits.

See Fig.3 for the front and rear panel's drilling diagram to suit the Ritec case.

Loading the firmware

Loading the firmware onto the Raspberry Pi Pico is quick and easy, thanks to the bootloader built into the Pico. All you need to do is hold down the white button on the top of the Pico while plugging its USB connector into your computer. The Raspberry Pi Pico will appear as a pseudo USB

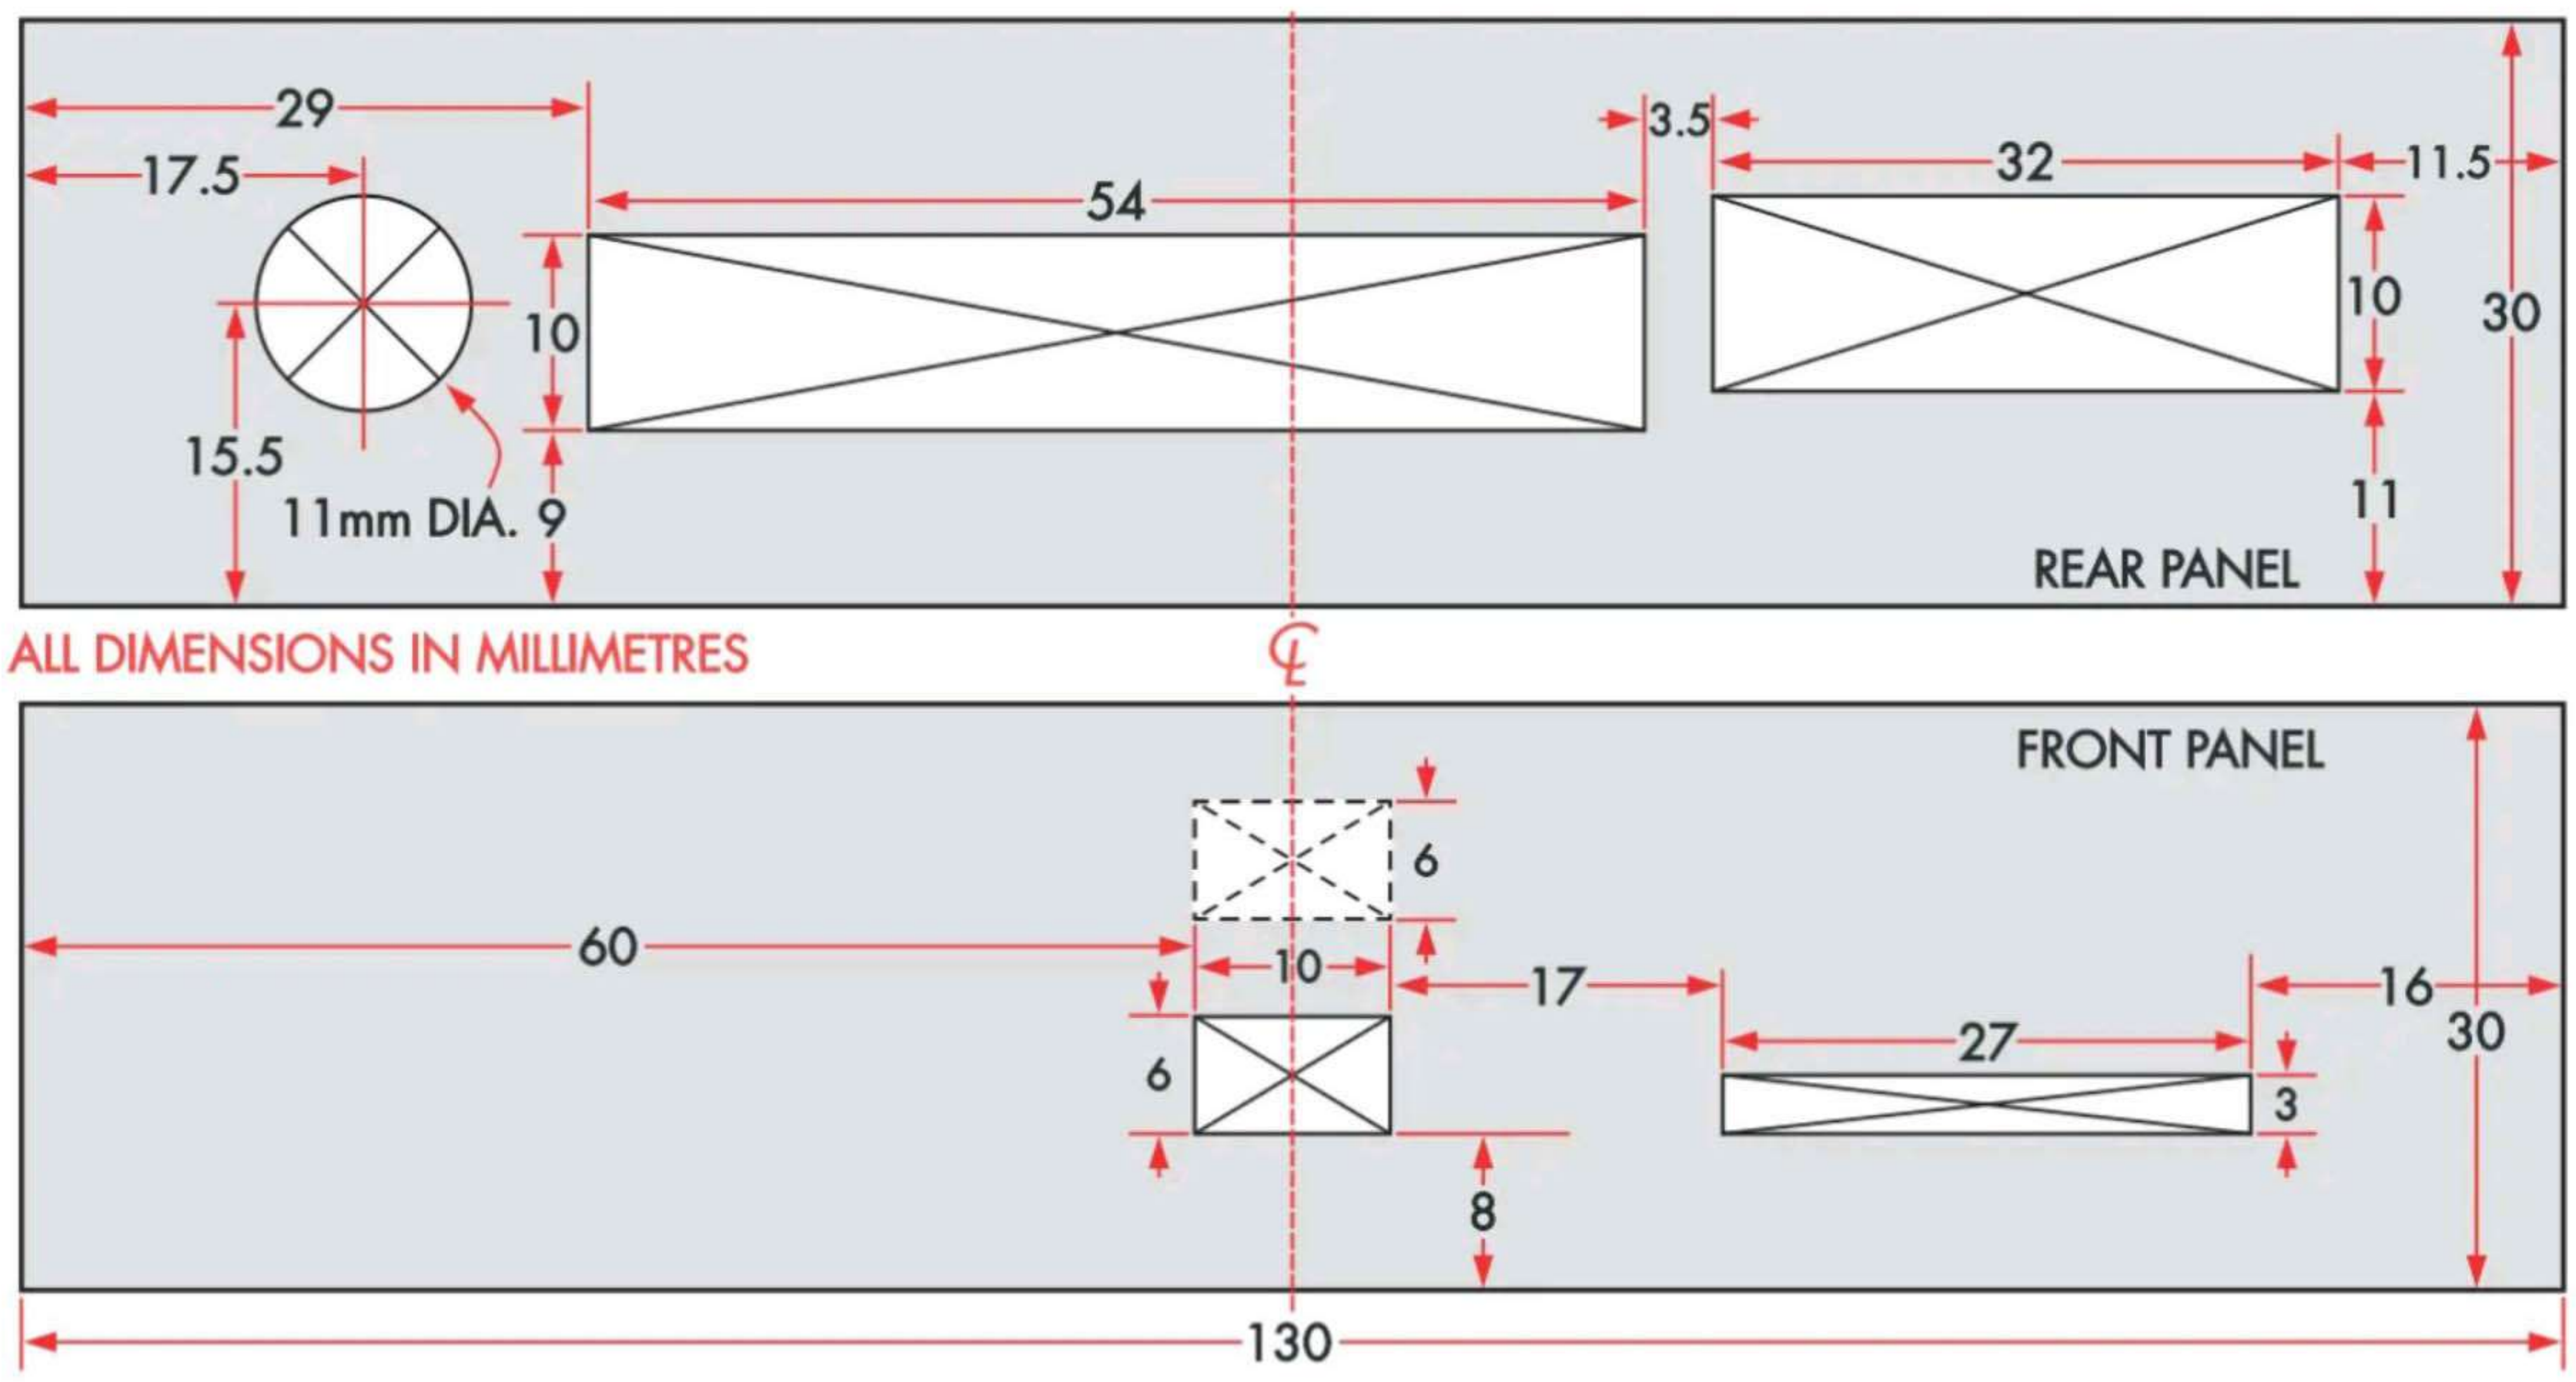

Fig.3: use these front & rear panel cutting diagrams/ templates to locate and size the holes in the case. The central cut-out in the front is to plug a cable into the USB socket on the Pico. Use the upper, dashed cut-out if you've mounted the Pico on headers. Either way, it might need enlarging depending on the size of the plug on your USB cable.
The distance between the dashed cutout on the front panel and the cutout below it is 5mm.

flash drive onto which you copy (eg, drag and drop) the "PicoMiteVGA.uf2" firmware file.

This will upload the file to the Pico and program it into its flash memory.

Following this, the Raspberry Pi Pico will drop the USB connection and reconnect as a virtual serial port over USB. But you can ignore that for the moment, as it will also come up driving the VGA output and looking for an attached keyboard.

If you have a VGA monitor attached, you will see the MMBasic copyright message (as shown in Screen 1) and, if you have a keyboard attached, you can try typing in a command. For example, try "PRINT 2 + 2" and, unless your computer failed kindergarten, it will display the number 4.

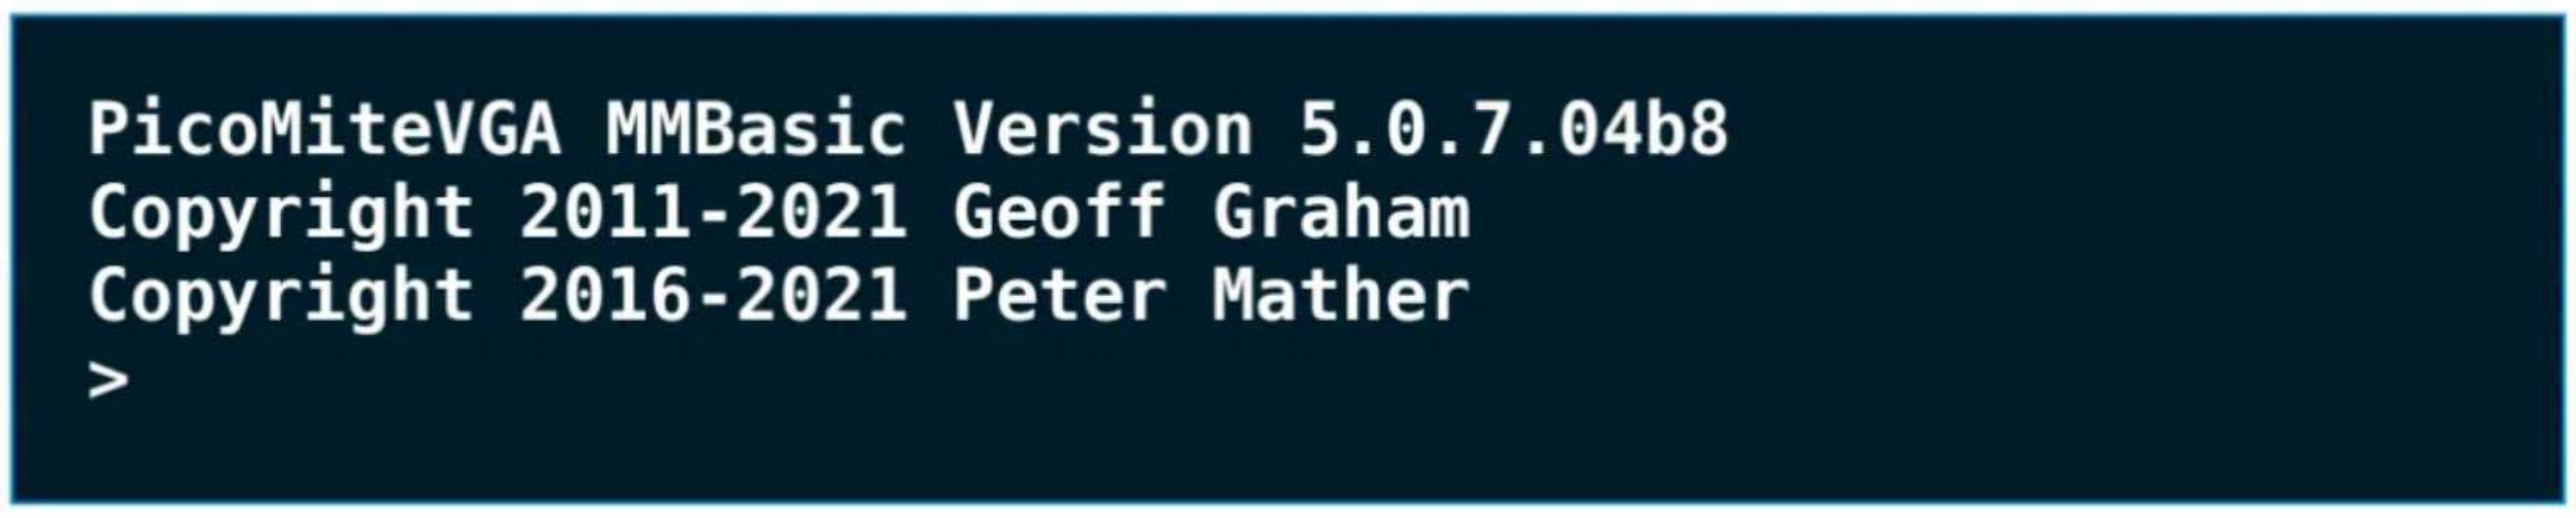

```
PicoMiteVGA MMBasic Version 5.0.7.04b8
Copyright 2011-2021 Geoff Graham
Copyright 2016-2021 Peter Mather
>
```

Screen 1: after loading the firmware, the VGA PicoMite will restart, activate the VGA output and look for an attached keyboard. With a VGA monitor connected, you will see this copyright message, and if you have a keyboard attached, you can try typing in a command or two.

The only adjustment needed is the 200Ω trimpot for the white balance. To correctly set this, you should display a large area of white. You can do this by entering the following line at the command prompt:

```
CLS RGB(white)
```

This will clear the screen to a white background, and you can then adjust the trimpot for a neutral white colour without a tint.

Fault finding

If you do not see anything on your monitor after loading the firmware, start by checking the green LED on the top of the Raspberry Pi Pico. It should be slowly flashing off/on with a period of about a second and a half. This indicates that the Pico is correctly programmed with the firmware and is running MMBasic.

If you do not see the LED flashing, reprogram the Raspberry Pi Pico and make sure that it completes successfully. Also make sure that you used the VGA PicoMite firmware and not the standard PicoMite firmware without the VGA capability – that will also flash the LED but not provide the VGA output.

If the Raspberry Pi Pico is OK, you will then need to resort to standard fault-finding procedures. Carefully check all solder joints with a magnifier, look for solder bridges, check all components and their values and make sure that the orientations of the diodes and Mosfets are correct.

It is also worth checking your monitor, keyboard and connecting cables. Do they work with other devices? The cables can be a problem; in the past, some constructors have been baffled by the lack of video only to discover that their VGA cable was faulty.

SD card

The VGA output and keyboard interfaces automatically operate after the firmware is installed, but the interface to the SD card needs to be configured before it can be used. This is done with the following command, entered on one line at the MMBasic command prompt:

```
OPTION SDCARD GP13, GP10,
  GP11, GP12
```

This tells the firmware that the SD card socket is connected and what I/O pins are used for its chip select ($\overline{CS}$), clock, MOSI and MISO functions, respectively.

Entering this option will cause the VGA PicoMite to reboot, but from then on, the setting will be remembered, even after power off. So you only have to enter it once.

After setting it up and inserting an SD card, you can test it by entering the command "FILES" at the prompt, and it should list the files and directories found on the card. If you had a BASIC program on the SD card, you could load it with the command:

```
LOAD "filename"
```

Then, if you edit it within MMBasic, you can save it back to the SD card with the command:

```
SAVE "filename"
```

The double quotes around the filename are required – you will get an error if they are not used.

As mentioned earlier, the SD card is not strictly necessary. The firmware reserves eight 'slots' in the Raspberry Pi Pico's flash memory for program storage on the chip.

So, if you have entered a program, you could save it into (for example) slot 6 with the command "FLASH SAVE 6". Later, you can load it back into program memory with the "FLASH LOAD 6" command.

What next?

To get a feel for the VGA PicoMite, you can enter the following short program. Start the process with the "EDIT" command. This is described in the user manual but, for the moment, all that you need to know is that in the editor, anything that you type will be inserted at the cursor, the arrow keys will move the cursor and backspace will delete the character before the cursor.

At the command prompt, type EDIT followed by the Enter key. The editor should start up, and you can type the following four lines:

```
DO
INPUT "What is your name?
  ", N$
PRINT "Hello, " N$
LOOP
```

Then press the F1 key on the keyboard. This tells the editor to save your program and exit to the command prompt. At the command prompt, type RUN and press the Enter key.

Your new computer should ask for your name, and when you type it in (followed by the Enter key), it will reply with a greeting. To break out of the program and return to the prompt, press Ctrl-C.

There you are; you have just written and run your first program on the VGA PicoMite. If you type EDIT again, you will be taken back into the editor, where you can change or add to your program.

You will probably have many questions at this point, and we have written a detailed user's manual to answer them. You can download this for free from the Silicon Chip website or the author's web page (https://geoffg.net/picomitevga.html). It should cover everything that you need to know.

If you are new to BASIC programming, open up this manual and check out Appendix G (Programming in BASIC – A Tutorial) near the end. This comprehensive tutorial on the language will take you through the fundamentals of programming in BASIC in an easy-to-read format with lots of examples. When you have finished, you will be a 'gung-ho' BASIC programmer!

Tetris

To test out this new computer, we decided to write a program to play the game of Tetris. Tetris is well-known and has been ported to over 65 platforms, a Guinness world record.

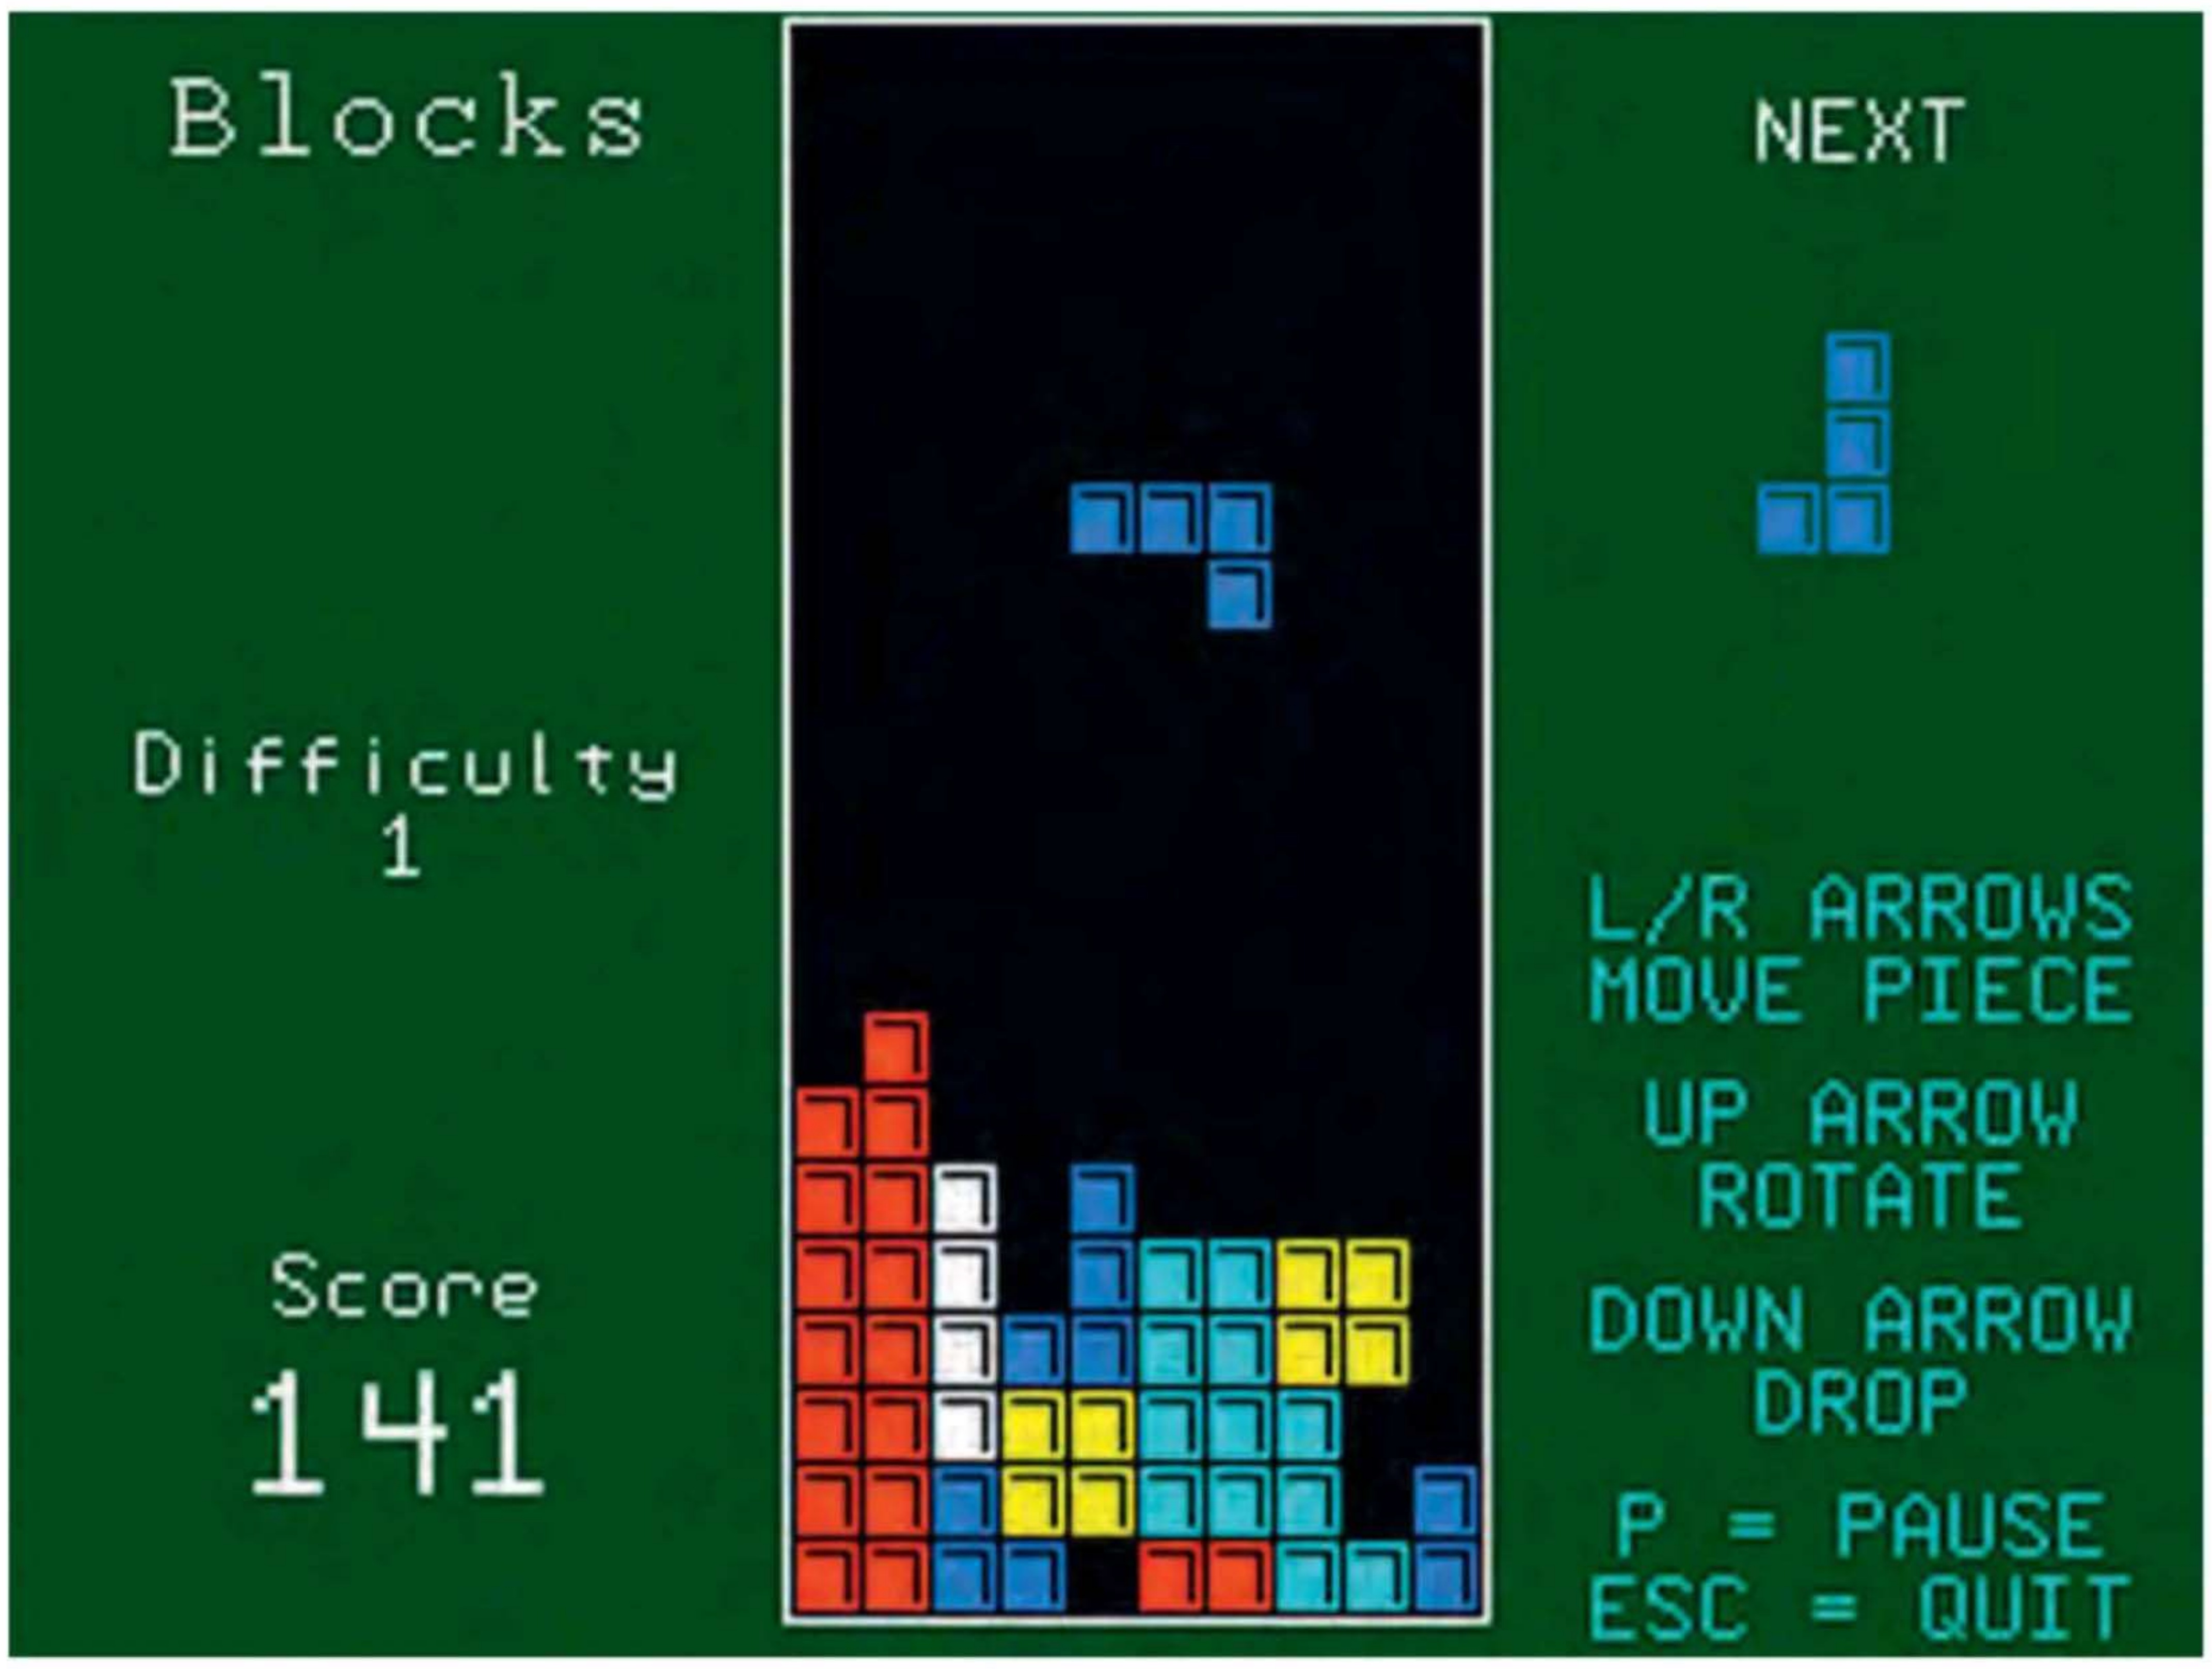

Screen 2: Blocks is a BASIC game that runs on the VGA PicoMite. It is colourful and fast, and it uses only 15% of the available program memory, so you can have fun adding to it if you wish. It will also run on the Colour Maximite 2.

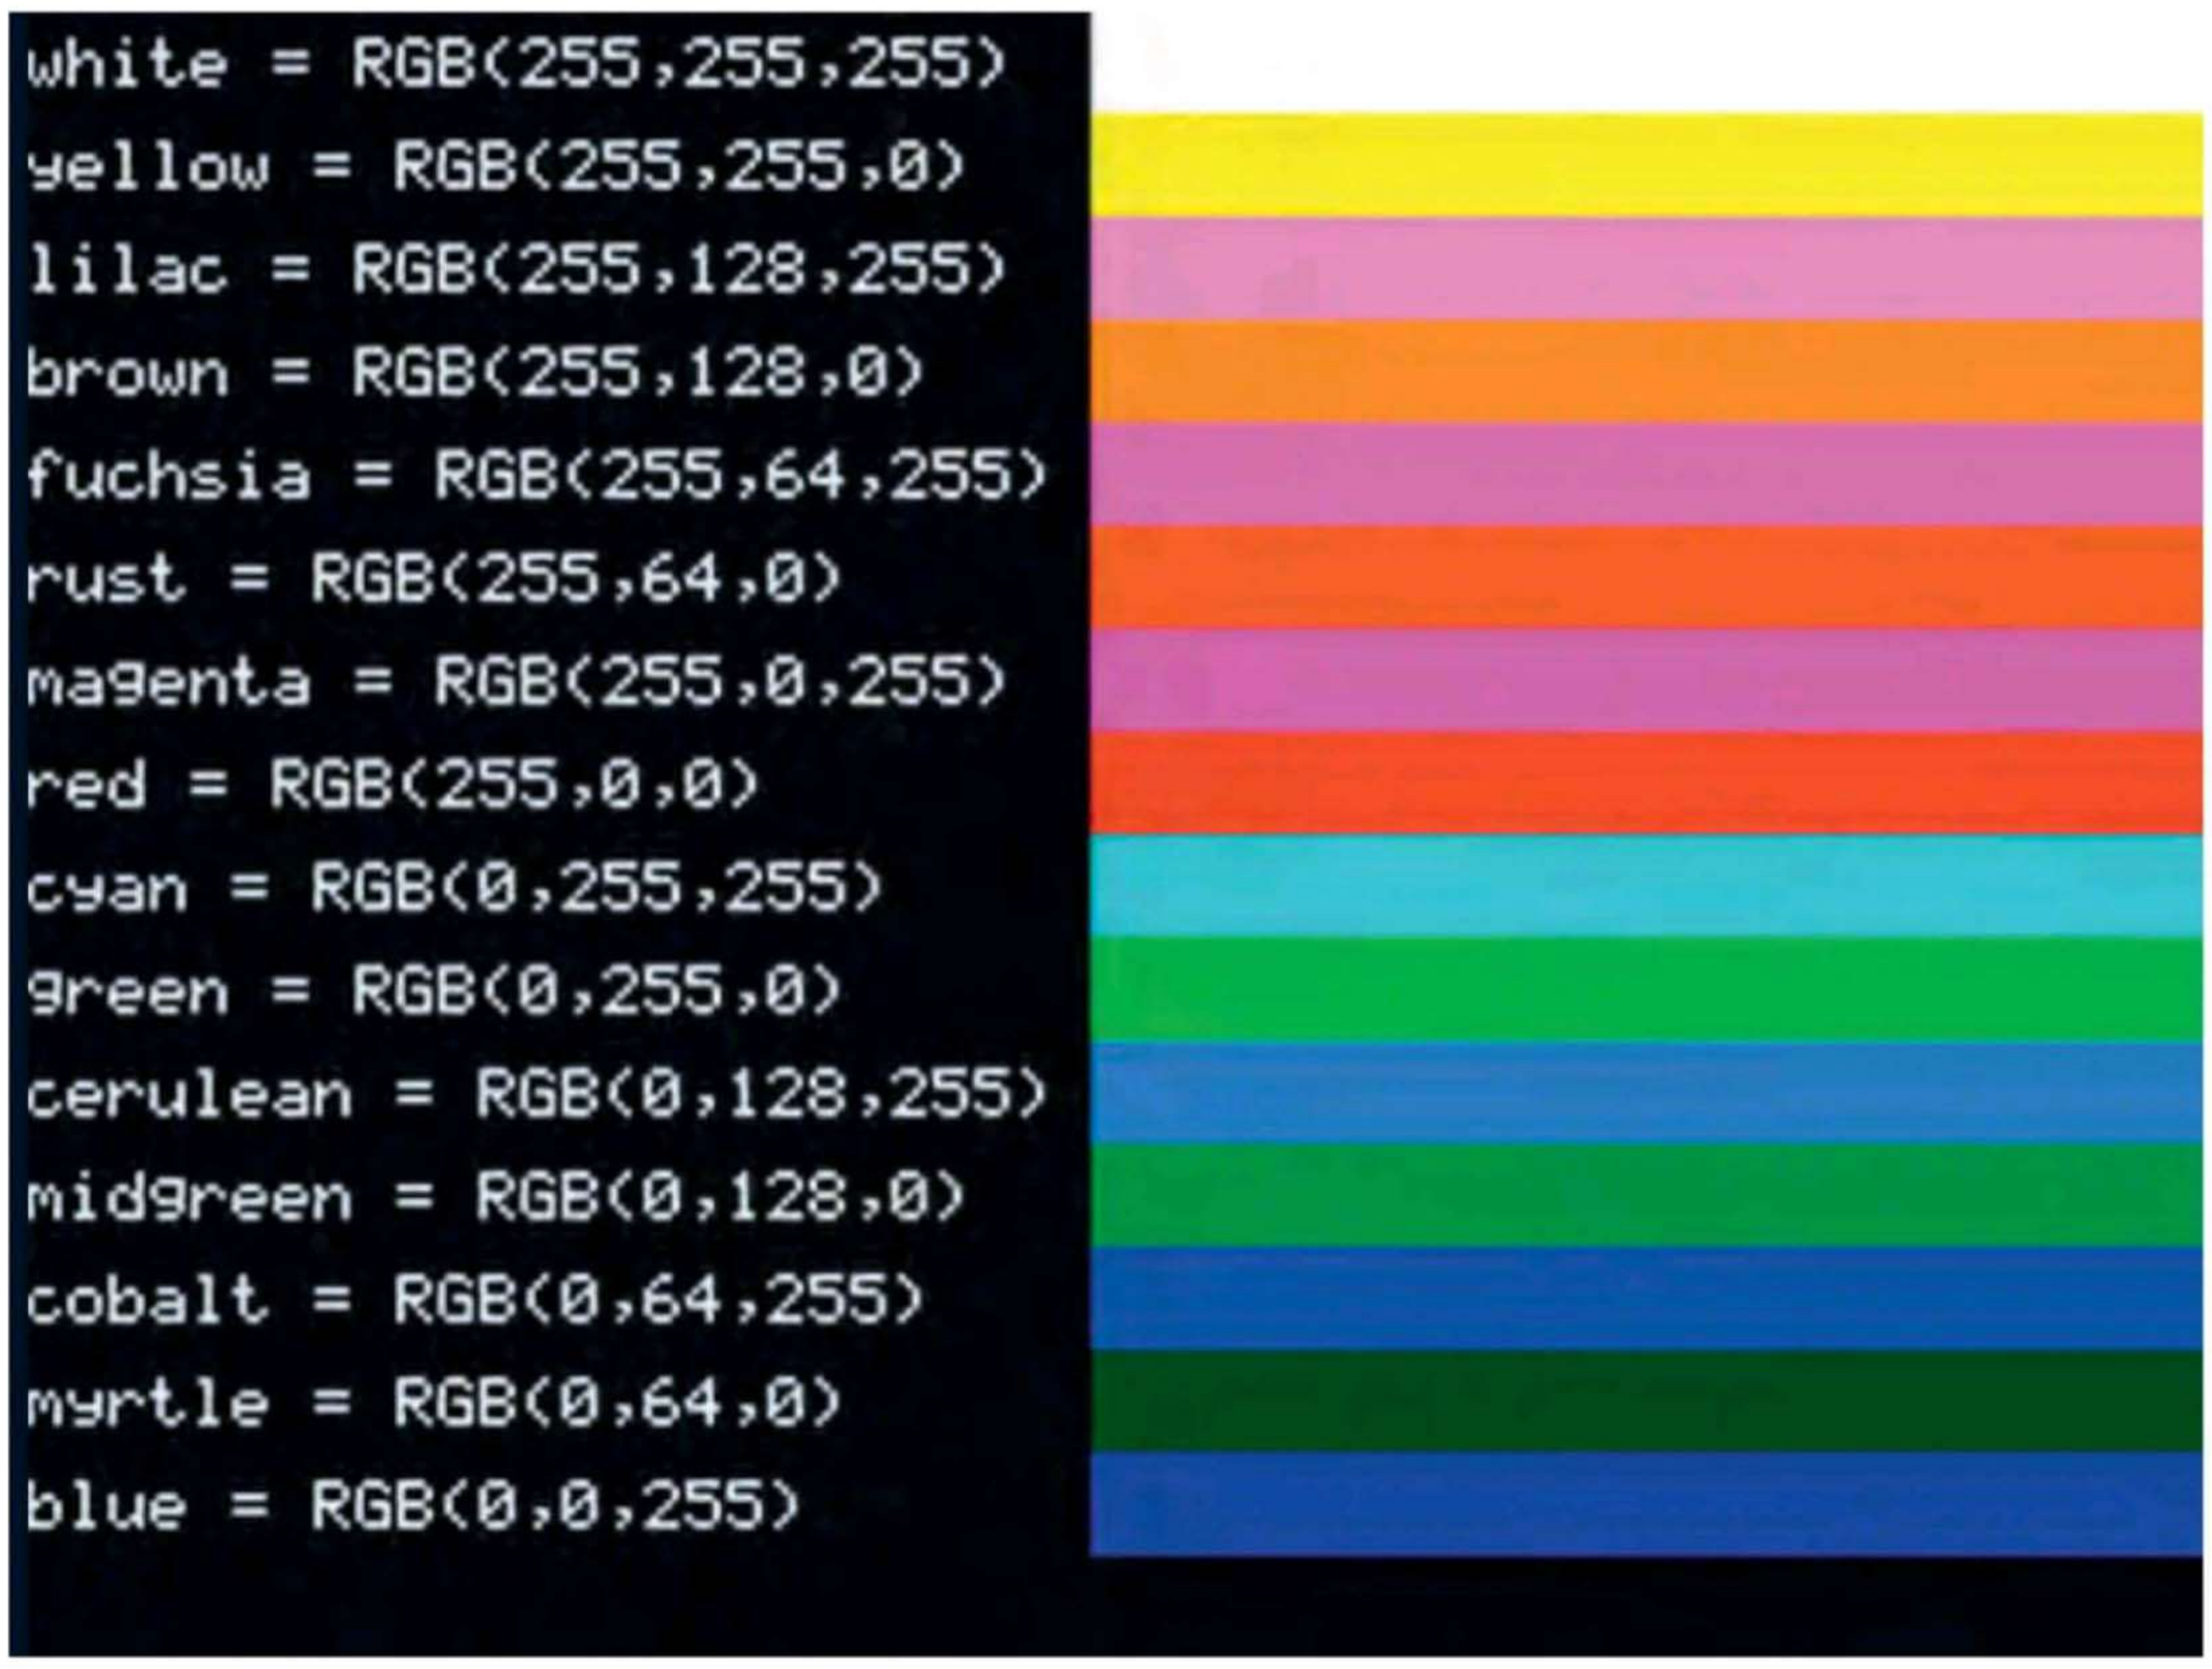

```
white = RGB(255,255,255)
yellow = RGB(255,255,0)
lilac = RGB(255,128,255)
brown = RGB(255,128,0)
fuchsia = RGB(255,64,255)
rust = RGB(255,64,0)
magenta = RGB(255,0,255)
red = RGB(255,0,0)
cyan = RGB(0,255,255)
green = RGB(0,255,0)
cerulean = RGB(0,128,255)
midgreen = RGB(0,128,0)
cobalt = RGB(0,64,255)
myrtle = RGB(0,64,0)
blue = RGB(0,0,255)
```

Screen 3: The output of the "Colours.bas" program, included with the VGA PicoMite firmware download. It shows all the colours the VGA PicoMite can generate. Run it on your computer to see the true colours as the printing press cannot reproduce all the colours accurately.

We were hoping to make it 66, but it turns out that the name and the game are copyrighted, so we developed something similar that we called Blocks, which we believe is just as much fun to play – see Screen 2. It is not a huge program (about 475 lines), and it uses just the basic graphics commands (line, box etc), so there is nothing special here.

Anyone could program this game with a bit of familiarity with BASIC, and you can easily dig into the code to see how it works.

On the VGA PicoMite, Blocks runs very fast, and it uses only 15% of the available program memory. This illustrates the capability of this little computer and underlines the fact that you can aim big with it, and it will not let you down.

The Blocks program is included in the VGA PicoMite firmware download so all you need to do is copy it to an SD card, then transfer that card to the VGA PicoMite and enter the command:

```
RUN "Blocks.bas"
```

Incidentally, Blocks will also run on the Colour Maximite 2, so if you have built that computer, you can try playing it on that.

There are two additional programs in the firmware download package. The first is "Fonts.bas", which will display the various fonts that come built into MMBasic on your monitor. This is handy when you are writing a program and need to select a suitable font.

The second program is called "Colours.bas" and it displays all 16 colours that the VGA PicoMite can generate on the monitor, including the codes to use in your BASIC program. Screen 3 shows what its output looks like, but the printing press will not accurately reproduce the colours, so run it on your VGA PicoMite to see the true set of colours.

So there you have it, one of the simplest 'boot to BASIC' computers possible. To keep up with firmware updates, including early 'beta' releases, check out the author's website at https://geoffg.net/picomitevga.html

It's also a good idea to visit the Back Shed forum (see www.thebackshed.com/forum/Microcontrollers) where there are many MMBasic fans swapping ideas and offering help to newcomers. PE

Using Low-cost Electronic Modules

By Jim Rowe

Self-Contained 6GHz Digital Attenuator

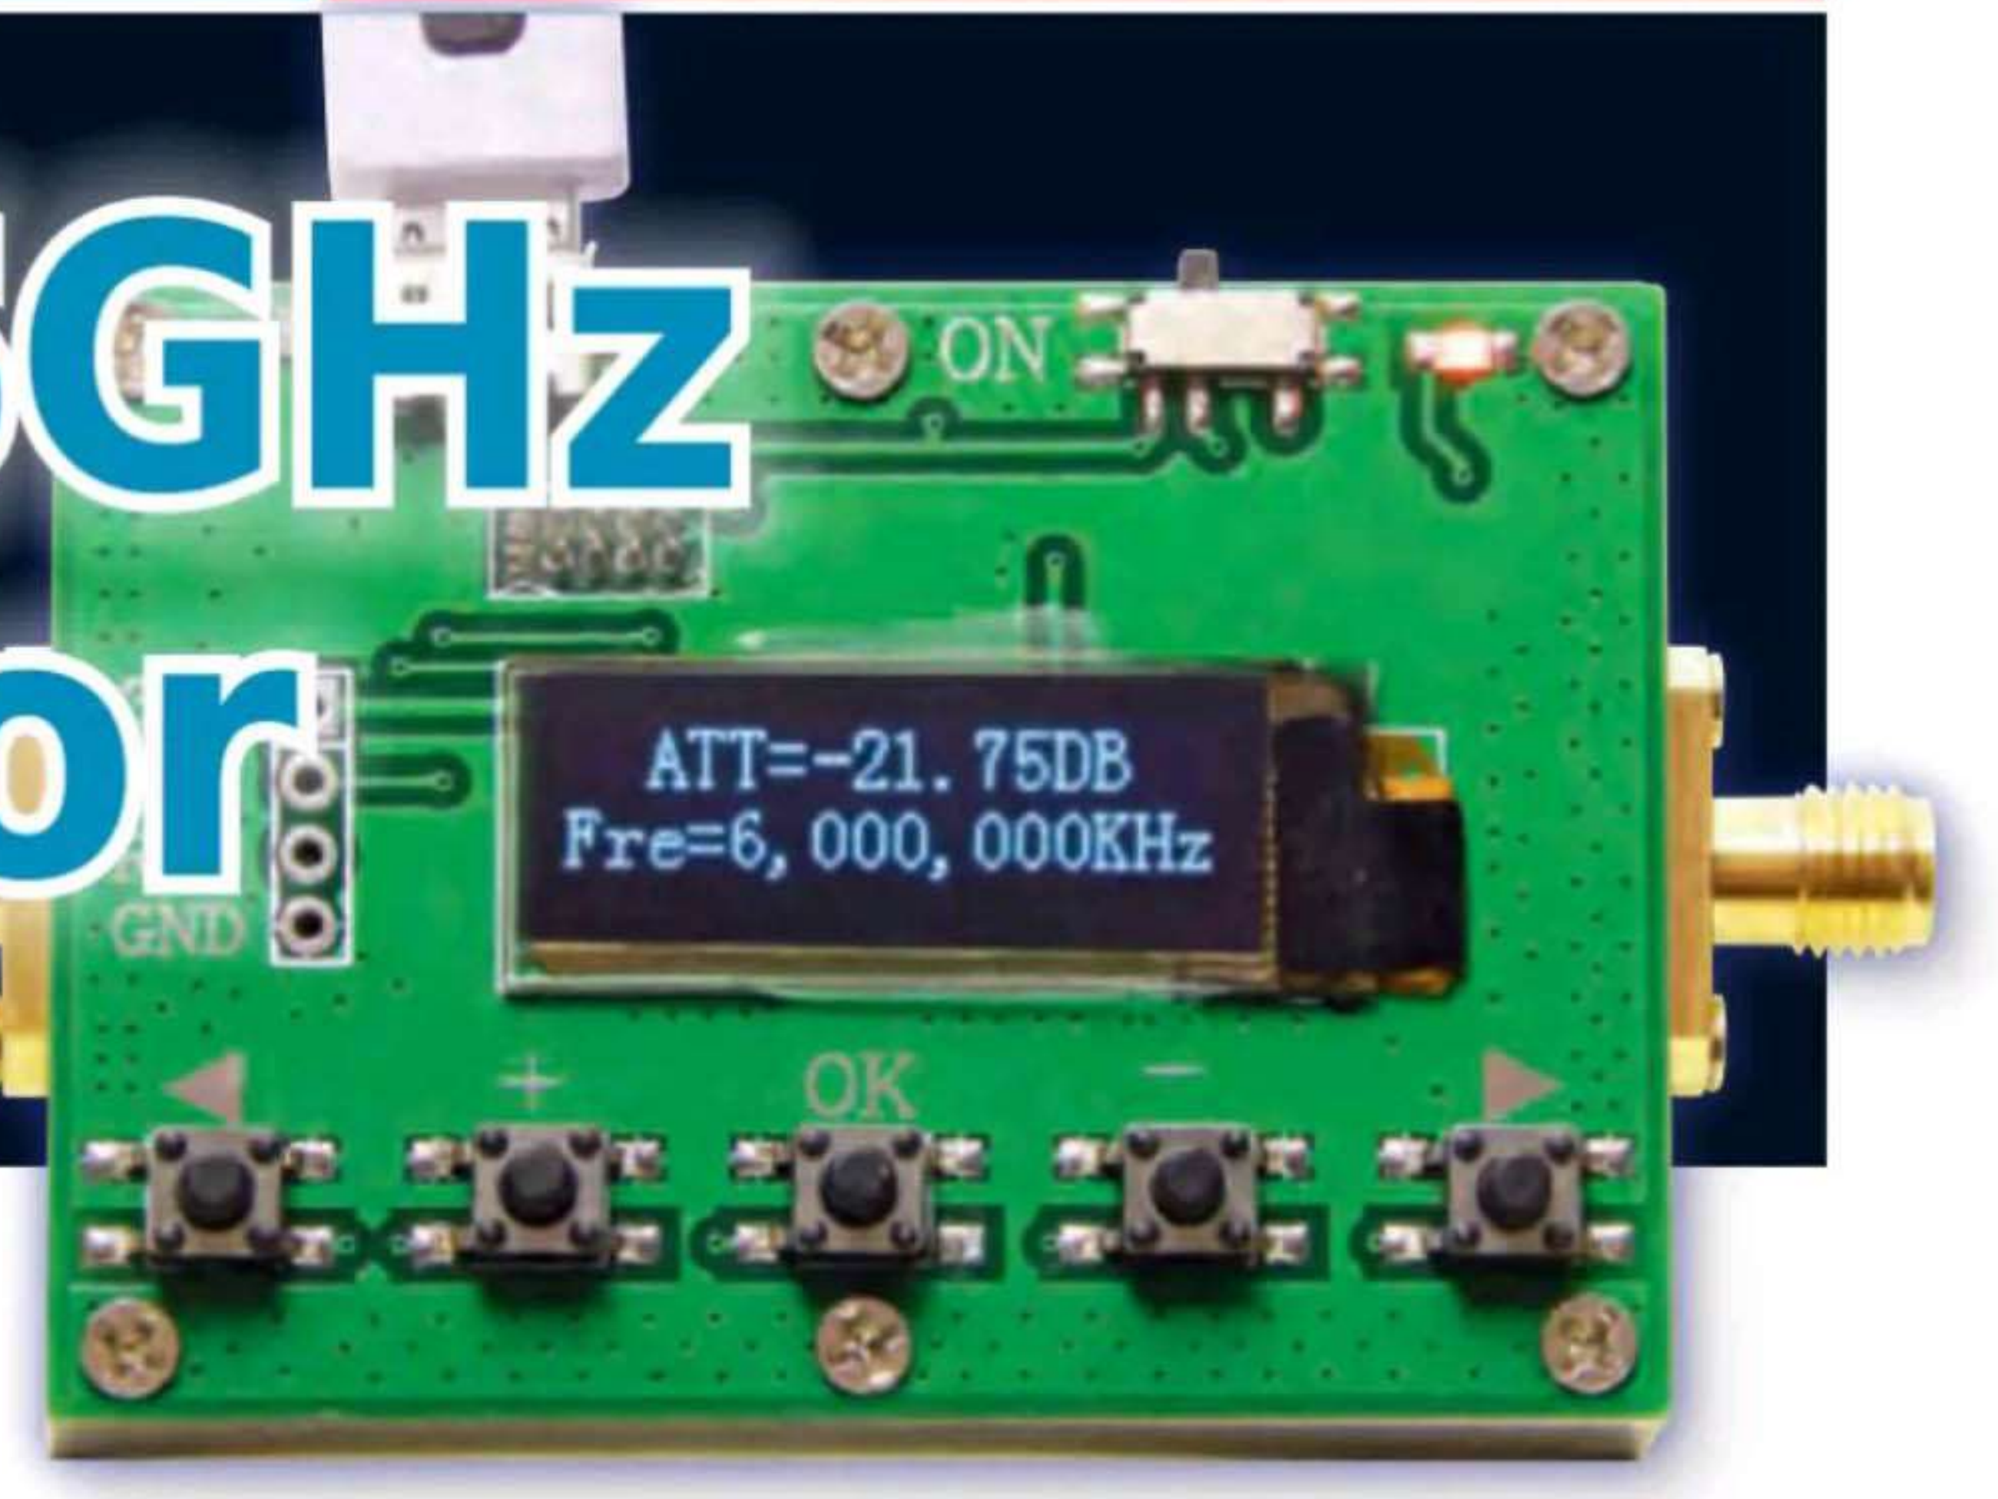

This digitally-programmable module can attenuate signals from 1MHz to 6GHz by 0 to 31.75dB in 0.25dB steps. You control it using five small pushbutton switches, while a tiny OLED screen shows the current setting.

I recently became aware of this digital attenuator that has a frequency range extending all the way to 6GHz, and 0.25dB attenuation steps from 0dB to 31.75dB. It seemed quite impressive, so I decided to review it.

It is available from a few different sellers online, such as the satisfyelectronics Store on AliExpress, at: https://vi.aliexpress.com/item/1005006358745427.html

You may be able to find the same module on other platforms like eBay and Amazon. If you want to buy the same one, make sure it matches the photos shown here. Some of the identifying features are the OLED screen in the middle of the module and the row of five tactile buttons along the bottom.

I paid AU$27 for it plus $9 postage for a total of $36 (about £18.50).

New module

The new module is likely available from several suppliers on the web, but I ordered the one shown in the photos from Banggood, catalog code 1648810. Currently, it's priced at $51.80 plus $6.70 for shipping to Australia. Like the earlier 3.8GHz module, it's almost certainly made in China.

The new module measures 56 x 40 x 16mm overall, not counting the SMA connectors at each end for RF input and output.

The digital attenuator section is on a small PCB fitted down inside a 56 x 40 x 10mm CNC machined aluminium block which forms the module's 'case'. The rest of the module's circuitry is mounted on a second PCB measuring 56 x 40mm, which forms the top of the case.

The OLED panel is mounted on the top of this PCB in the centre, along with the micro-USB power socket, the mini slider power switch and a tiny SMD power LED. Then along the PCB front are the five small pushbutton switches used to select the attenuation setting. Presumably, the rest of the controller circuitry is mounted on the underside of this PCB.

The UHF attenuator chip is probably the Analog Devices HMC1119, a 'big brother' to the HMC472 used in the aforementioned 3.8GHz attenuator.

According to the Analog Devices data sheet, the HMC1119 has a range of 100MHz to 6.0GHz and seven control bits, giving a setting range of 0 to 31.75dB in 0.25dB steps. It has a specified insertion loss of 1.3dB at 2.0GHz, drooping to around 1.5dB at 3.5GHz and a whisker below 2.0dB at 6GHz. Pretty impressive!

As with the 3.8GHz attenuator, I couldn't find a full circuit for the new module, so I could only work out a basic block diagram for it, shown in Fig.1.

The RF1 input and RF2 output pins of the HMC1119 chip are coupled to the SMA input and output connectors via capacitors. Apart from various bypass capacitors, that makes up all of the actual attenuator section.

Below is the control section, based on a microcontroller (possibly an STM32F103C8T6, like the one used in the 3.8GHz attenuator).

Operation

The microcontroller (MCU) controls the attenuation settings of the HMC1119 via the seven programming lines, while the user determines the attenuation setting using the five small pushbutton switches S1-S5. To make this easy, the MCU displays the current attenuation setting on the OLED screen, controlled using a standard I^{2}C serial interface.

When power is first applied, the MCU sets the attenuation to 00.00dB.

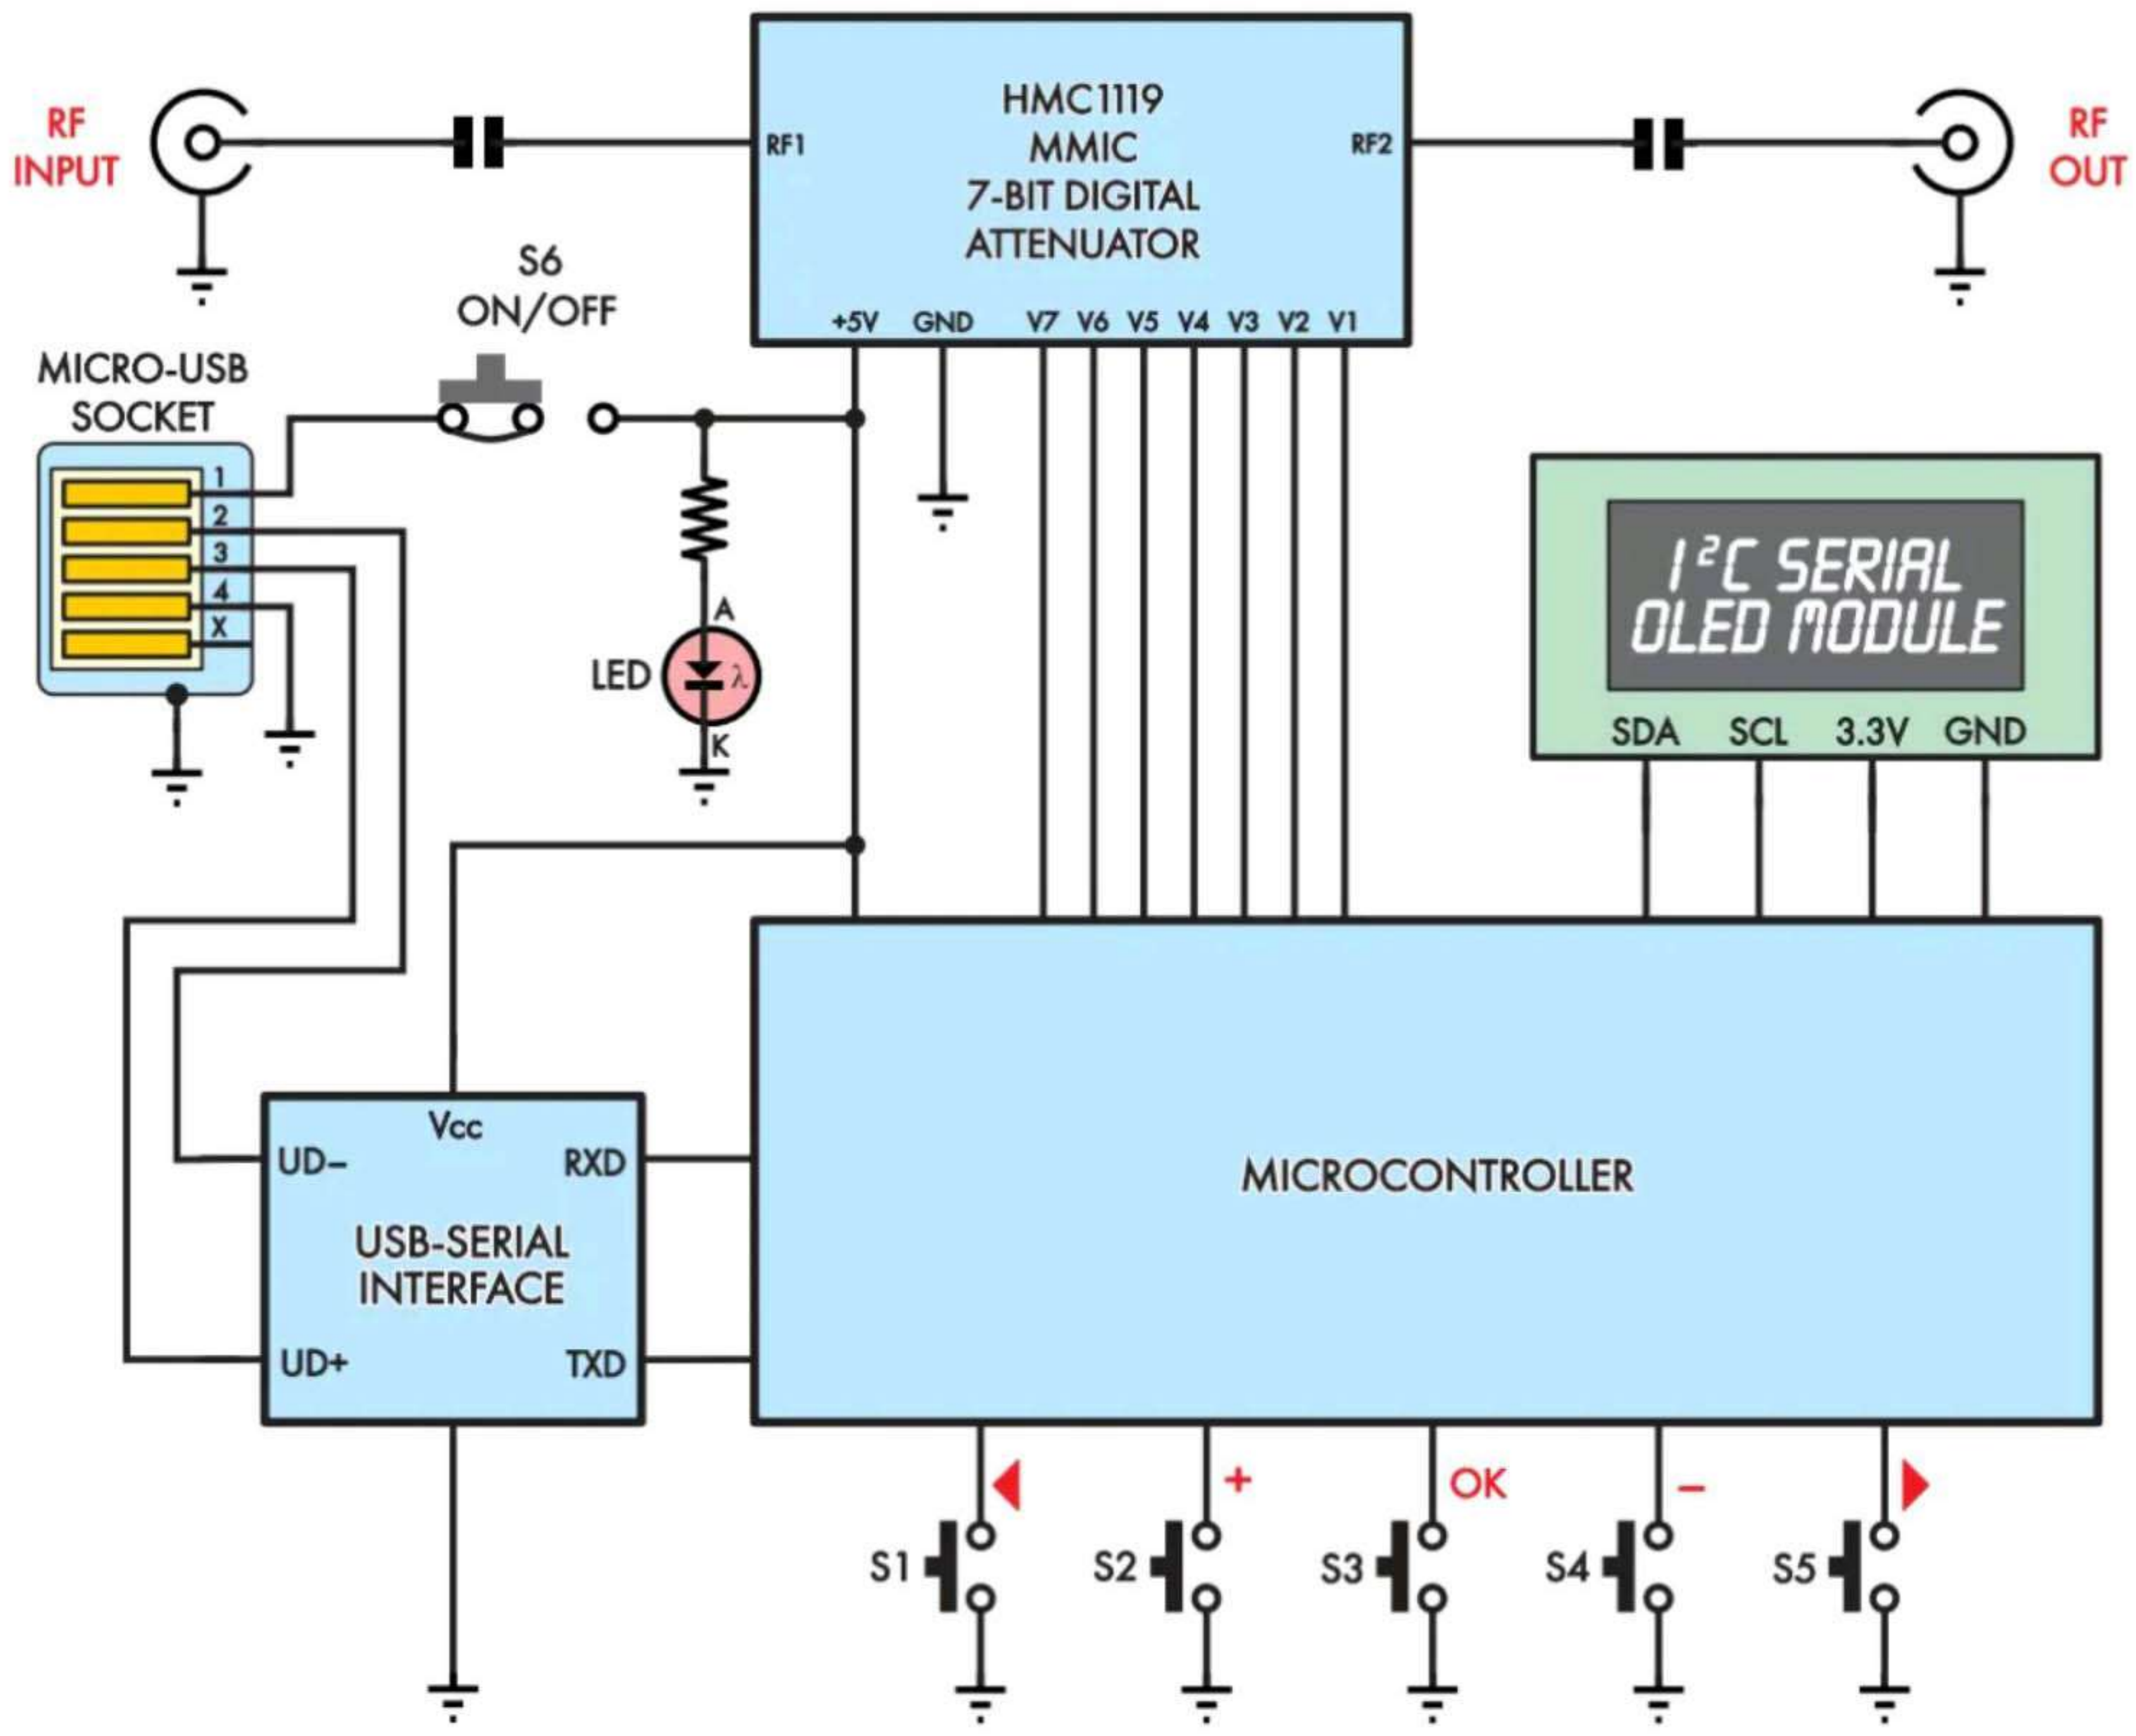

Fig.1: a simplified version of what we expect the block diagram the 6GHz attenuator to look like, as there is no full circuit diagram available.

To change this, you first press S3 (the OK button) and then press S1 (<) or S5 (>) until the display is flashing the setting digit you want to change. Then you can press either S2 (+) or S4 (-) to change the value of this digit. To change other digits, use either S1 or S5 to move to them, then use S2 or S4 to change their value. Then if you press S3 again, this will be the new setting.

It's pretty straightforward, and although the tiny pushbuttons used for S1-S5 seem to be the same as those used on the 3.8GHz module, the additional two buttons seem to allow the setting to be changed more reliably. Perhaps the firmware in the MCU has also been improved to make it less susceptible to contact bounce.

I have also shown a USB-serial interface chip in Fig.1. This chip may or may not be in the 6GHz module; I've shown it purely because it was present in the 3.8GHz module. It's possible that, in this case, the data lines from the micro-USB connector go directly to two pins of the MCU, but they certainly are routed somewhere on the PCB.

Either way, it would allow the attenuation setting to be programmed from an external PC, as well as from its own 'keyboard'. So the micro-USB socket is not just to feed power to the module, but also for external control.

As with the 3.8GHz attenuator, there's virtually no information provided on doing this external control, but I found a very cryptic suggestion in the 'Customer Q&As' section of the Banggood info on the module: "Go to github.com/emptemp/att6000_control for Python code."

I'm not familiar with the Python programming language, so I sought help from other Silicon Chip staff. They advised me that all the 'att6000' Python code seemed to do was send serial text commands in the format "wv0XXYY<LF>", where the XXYY characters indicate the desired attenuation setting XX.YY.

In other words, sending the command "wv02375<LF>" should change the attenuator's setting to -23.75dB.

They also informed me that the command should be sent at 115,200 baud, not the 9600 baud that seemed to be used previously. I did try this out, and the results are described below.

Performance

I measured the performance of the new attenuator module using my Signal Hound USB-SA44B HF-UHF spectrum analyser and its matching USB-TG44A tracking generator. Both were controlled by Signal Hound's Spike software (V3.5.15) in its SNA (scalar network analysis) mode.

Since the SA44B and TG44A combination will only work up to 4.4GHz, I could only check the module over this range.

ATT=-21.75DB
Fre=6,000,000KHz
3.3V
TXD
RXD
GND
ON
ATT-6000-30DB

The 6GHz digital attenuator from Banggood has an OLED screen and weighs about 57g.

I first used this setup to check the module's performance at an attenuation setting of 00.00dB to see its insertion loss. This is shown in Fig.2; the measured insertion loss is less than -2.5dB up to about 1.3GHz, then droops down to about -6.0dB at 2.5GHz, then improves to about -2.5dB at 3.0GHz.

It then droops to about -4.5dB at 4.0GHz, before moving up again to reach -4.0dB at 4.4GHz, which looks promising for its insertion loss at frequencies up to 6GHz.

After this, I did response tests at 'major' attenuation steps: -5dB, -10dB, -15dB, -20dB, -25dB and -30dB. These settings were chosen to give a good idea of the module's overall performance. After examining the results I then checked the response at a number of 'fine detail' settings: -1dB, -1.5dB, -2dB, -3dB, -4dB, -7.5dB, -10.75dB, -14dB, -19dB, -28.25dB and -31.75dB.

During each of these tests, I saved an image of Spike's plot of the test results. Then, knowing that there wouldn't be enough space to reproduce all 18 of the plots separately, I combined all of the plots into a single composite plot, to allow for easier evaluation – see Fig.3.

The upper plots in Fig.3 (down to about -20dB) have a shape almost identical to that of the top 00.00dB plot, just separated from it by the chosen attenuation setting.

For frequencies above about 1.75GHz, the higher attenuation plots (-20dB and greater) develop an increasing number of bumps and dips. These are very apparent in, for example, the red -25dB plot, the purple -28.25dB plot, the red -30dB plot and the blue -31.75dB plot.

All of these four plots show an increasing tendency to have a significant dip between 2.5GHz and 3.1GHz. I suspect that this may be due to small resonances inside the HMC1119 chip and/or its surrounding tracks on the attenuator section's PCB. There might also be standing waves inside the attenuator box at specific frequencies.

These plots tell us that the attenuator's performance is quite respectable, at least for frequencies up to about 2.2GHz and for settings up to about -20dB. But the errors do increase for frequencies above 2.2GHz and with settings above -20dB.

Of course, the attenuator would still have many practical uses at frequencies above 2.2GHz and with settings above -20dB, especially if you were to use Fig.3 to correct for likely errors.

Armed with the information mentioned earlier on how to control the device over a serial connection, it didn't take me long at all to test sending new attenuation settings from my Windows 10 PC, using the TeraTerm serial terminal application.

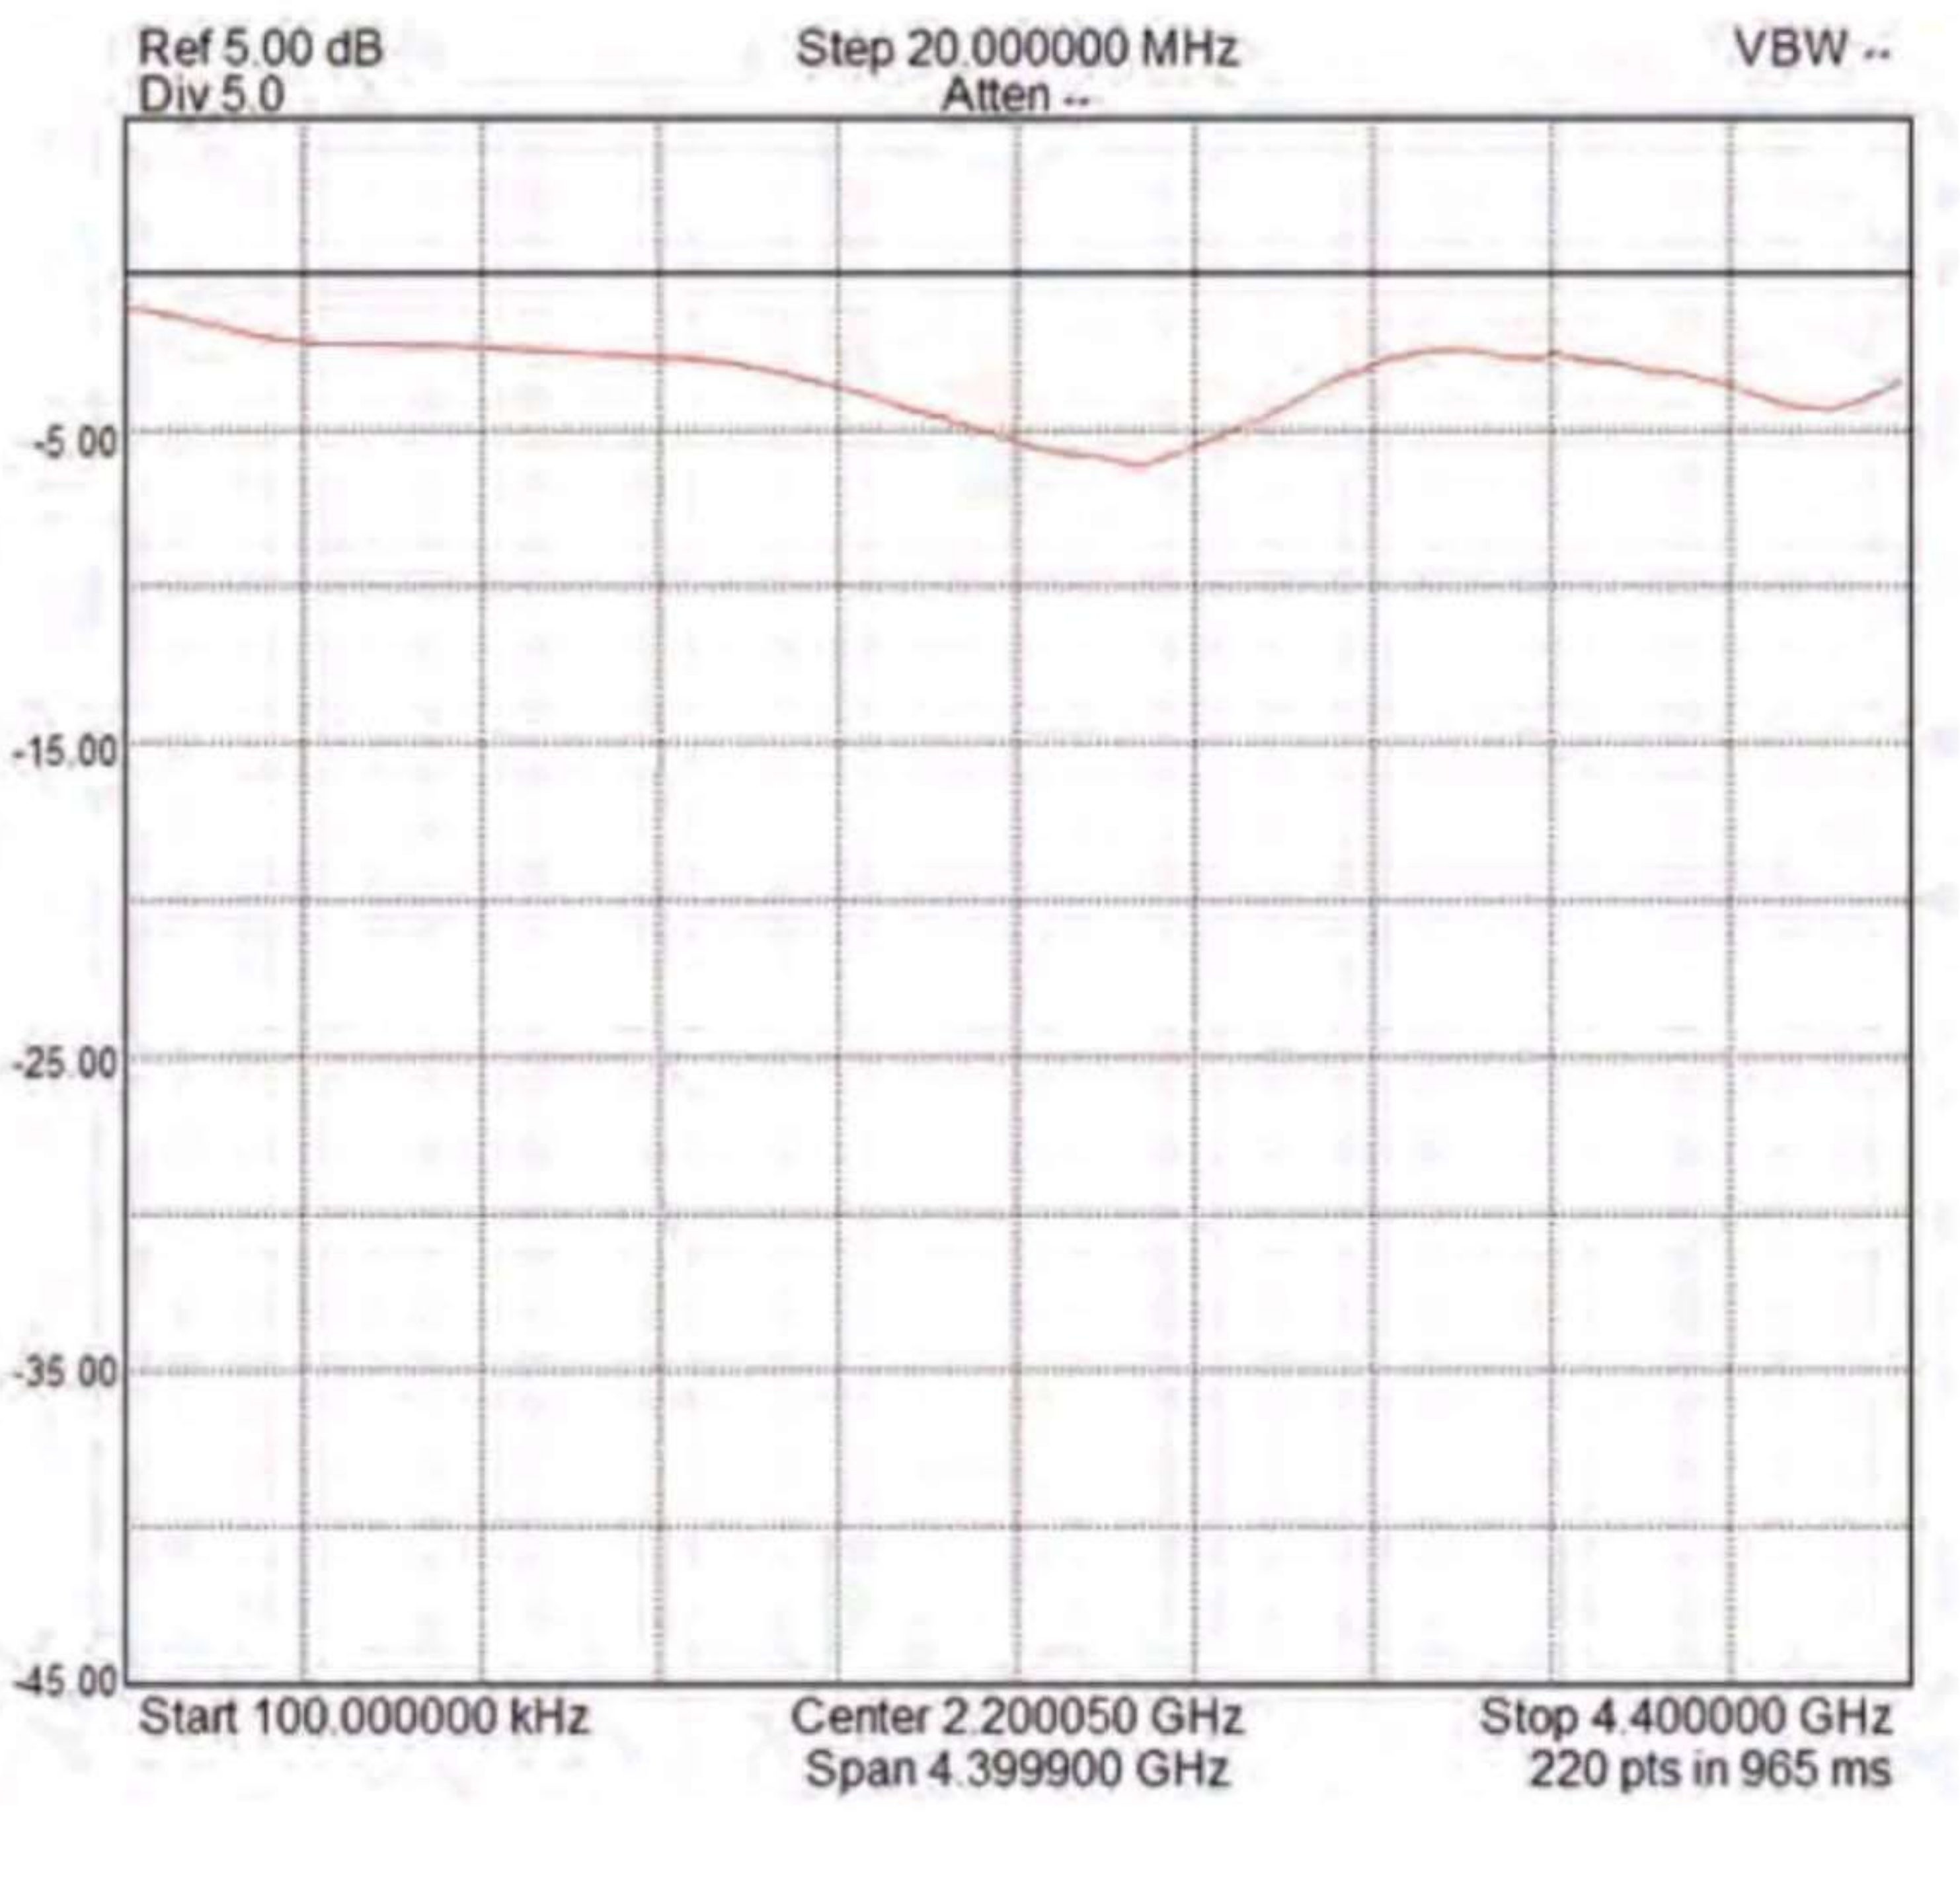

Fig.2: using Signal Hound's Spike software the 6GHz module could be checked at an attenuation setting of 0dB to measure its insertion loss. Note that the setup used for testing can only measure up to 4.4GHz, so not the full range of the attenuator.

All I had to do was plug the cable from the attenuator into a USB port, then go into Settings → Devices to find out to which Virtual COM port it had been assigned. Then I started up Tera-Term and set it up to communicate with that port at 115,200 baud, with the 8N1 data format and with only an LF (line feed) at the end of each line.

I was then able to change the attenuator's setting at any time simply by typing in a command like "wv01575" and pressing the Enter key. No problem! The attenuator's OLED immediately showed the new setting (like "-15.75dB") and also sent back an "OK" message, to confirm that the command had been received and acted upon.

I should perhaps note that there does seem to be provision on the top of the attenuator (just to the left of the OLED) for connecting a separate serial interface, as you can see in the photos. But there's no information on doing this. I guess that the command interface is the same, but I haven't tried it, so I can't say for sure.

Conclusions

Overall this new attenuator module seems reasonably good value for money when you consider its relatively wide frequency range and low price. I also like its ability to be programmed using the built-in MCU, control buttons and tiny OLED screen, or from a PC via the USB port (and presumably from a separate microcontroller, via the serial port header).

My only real gripe is that when I tried to unplug the USB cable from the micro-USB socket after testing it, the socket lifted straight off the PCB. It seemed to have been poorly soldered, and as a result, I had to spend quite a bit of time soldering it back on (under a microscope). I'd have preferred a mini-USB socket, as these seem to be a bit more rugged and also attach more securely to the PCB. **PE**

ENIGMA
COMPONENT SHOP
www.enigma-shop.com
Electronics Components for the Hobbyist!
Cable & Wire • Capacitors • Connectors • Filters
Kits & Modules • Hardware & Cooling
and much more...
READY TO BE REWARDED?
Join our loyalty scheme and **earn points** on every purchase you make which can be **used for discounts** in the future.
Read our reviews on Trustpilot

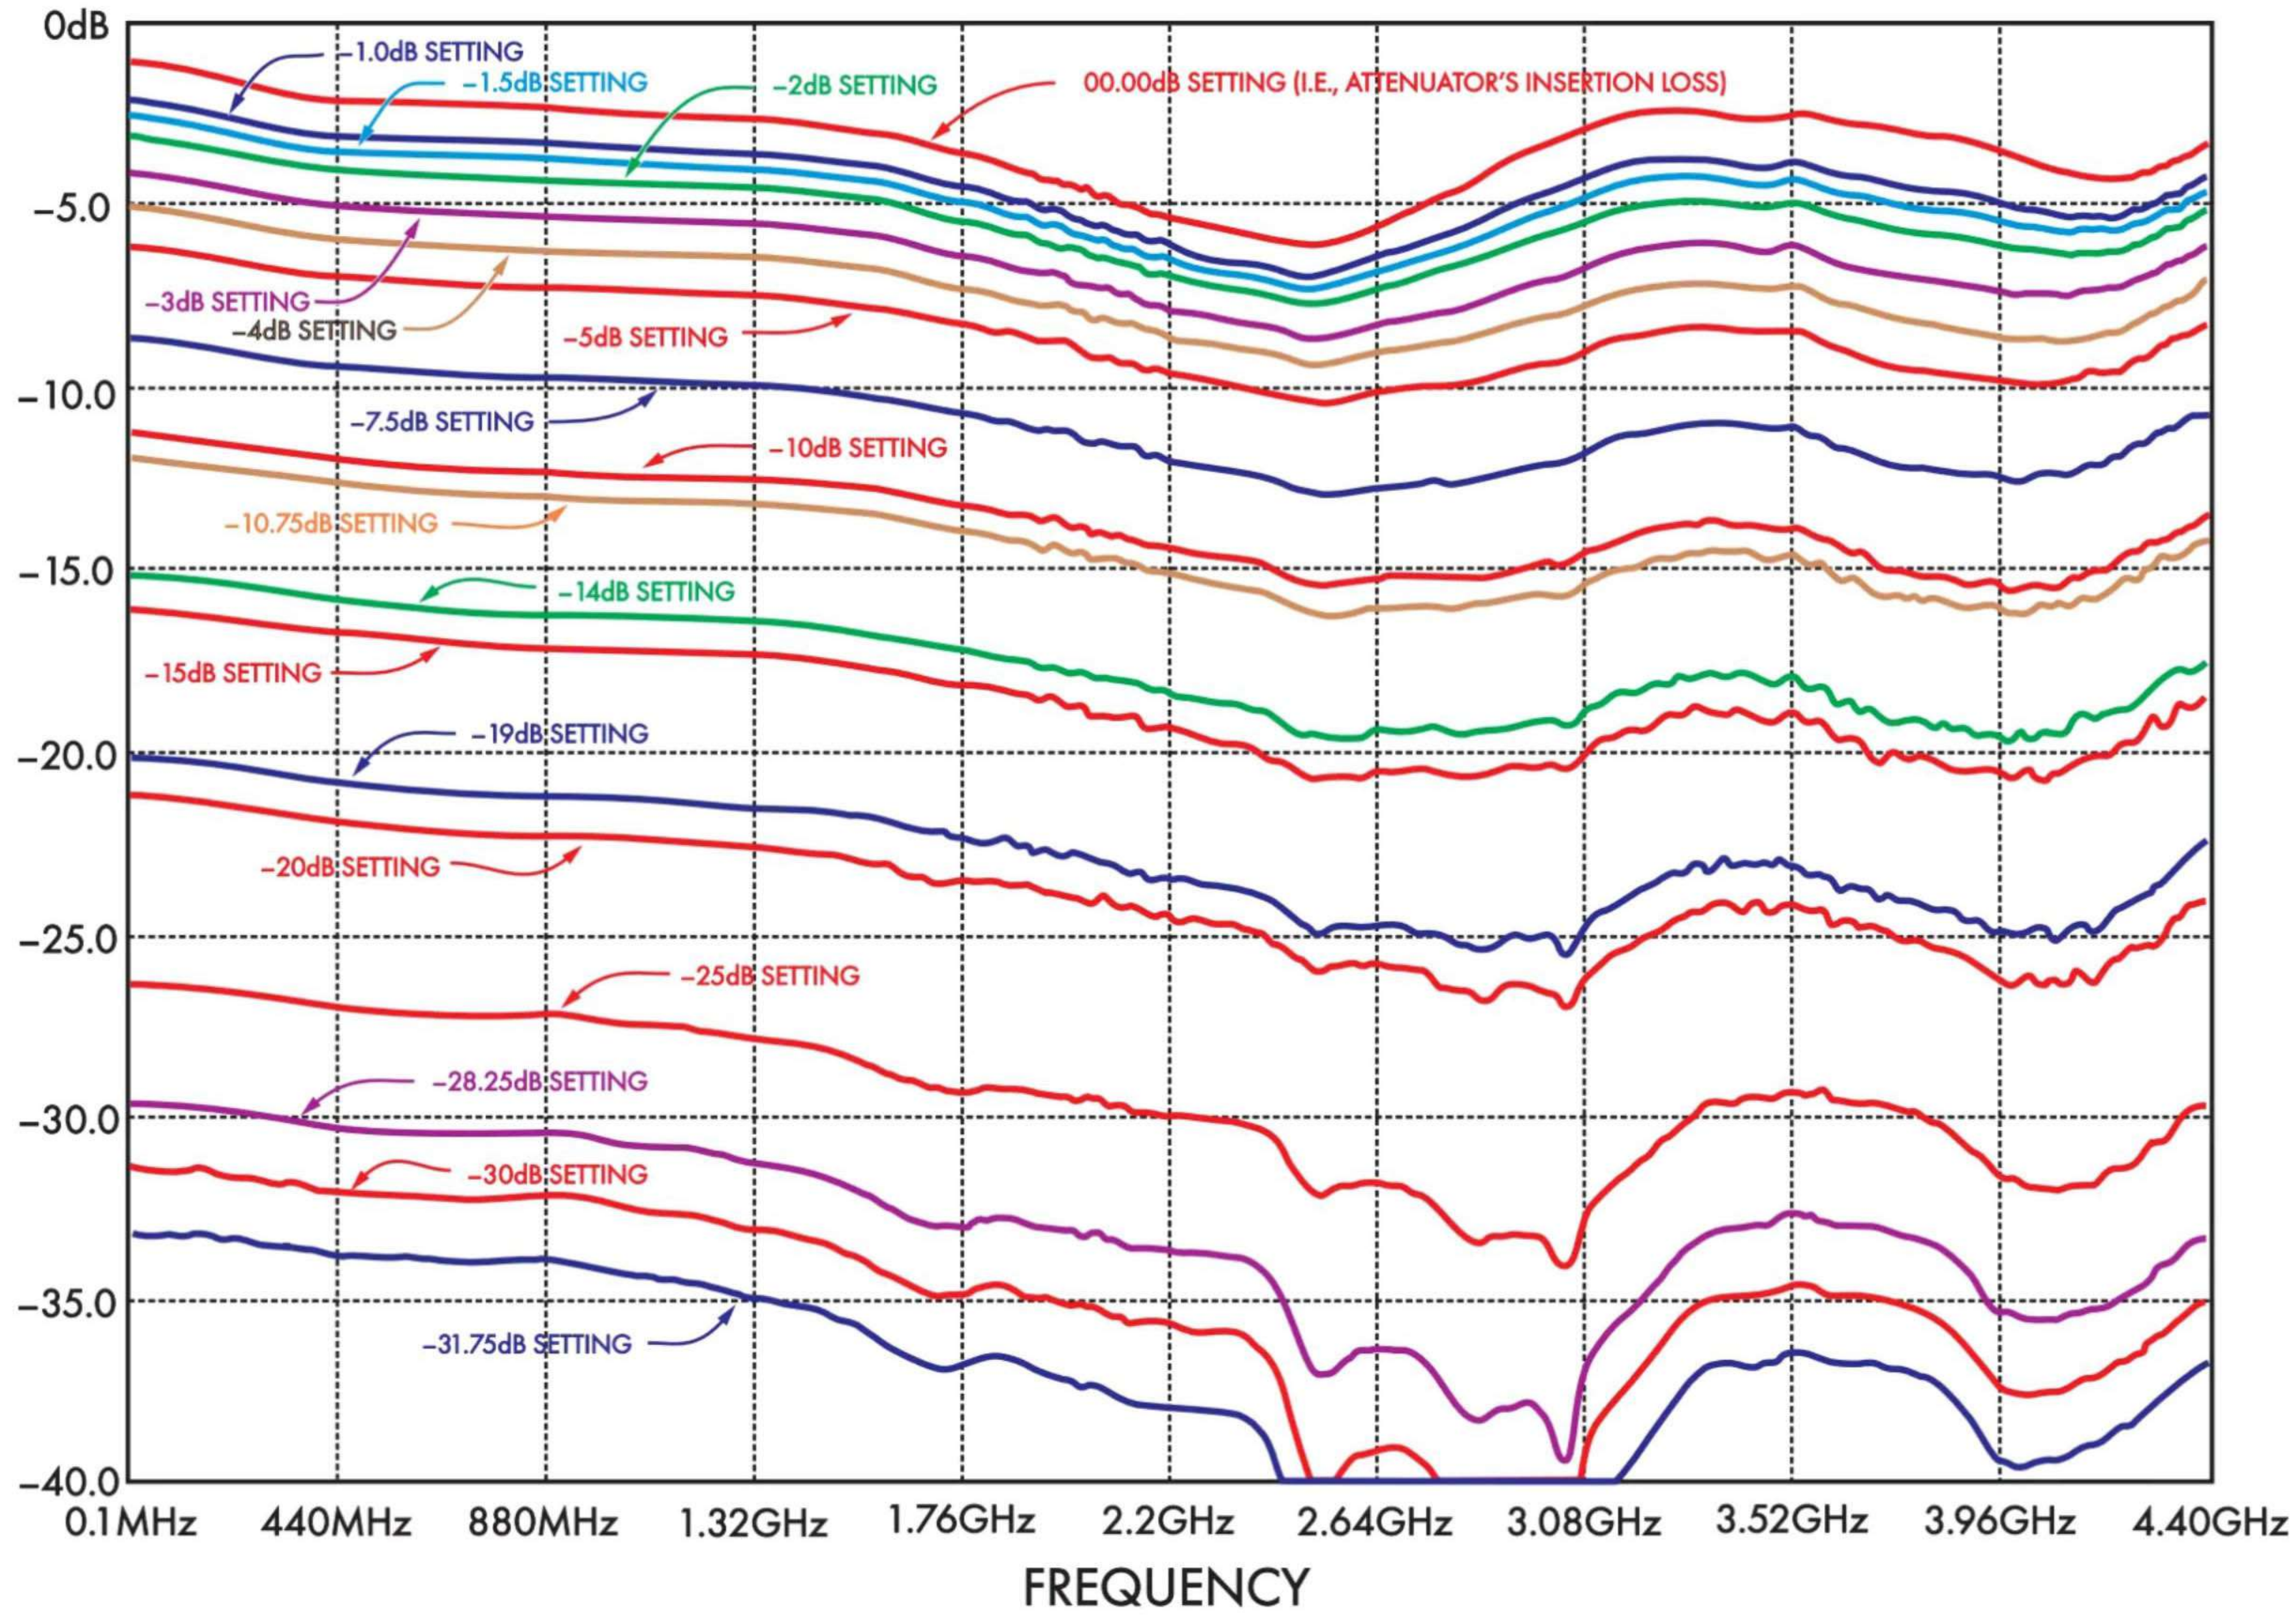

Fig.3: a graph showing the combined result from a variety of response tests on the attenuator at various settings.

Feature article

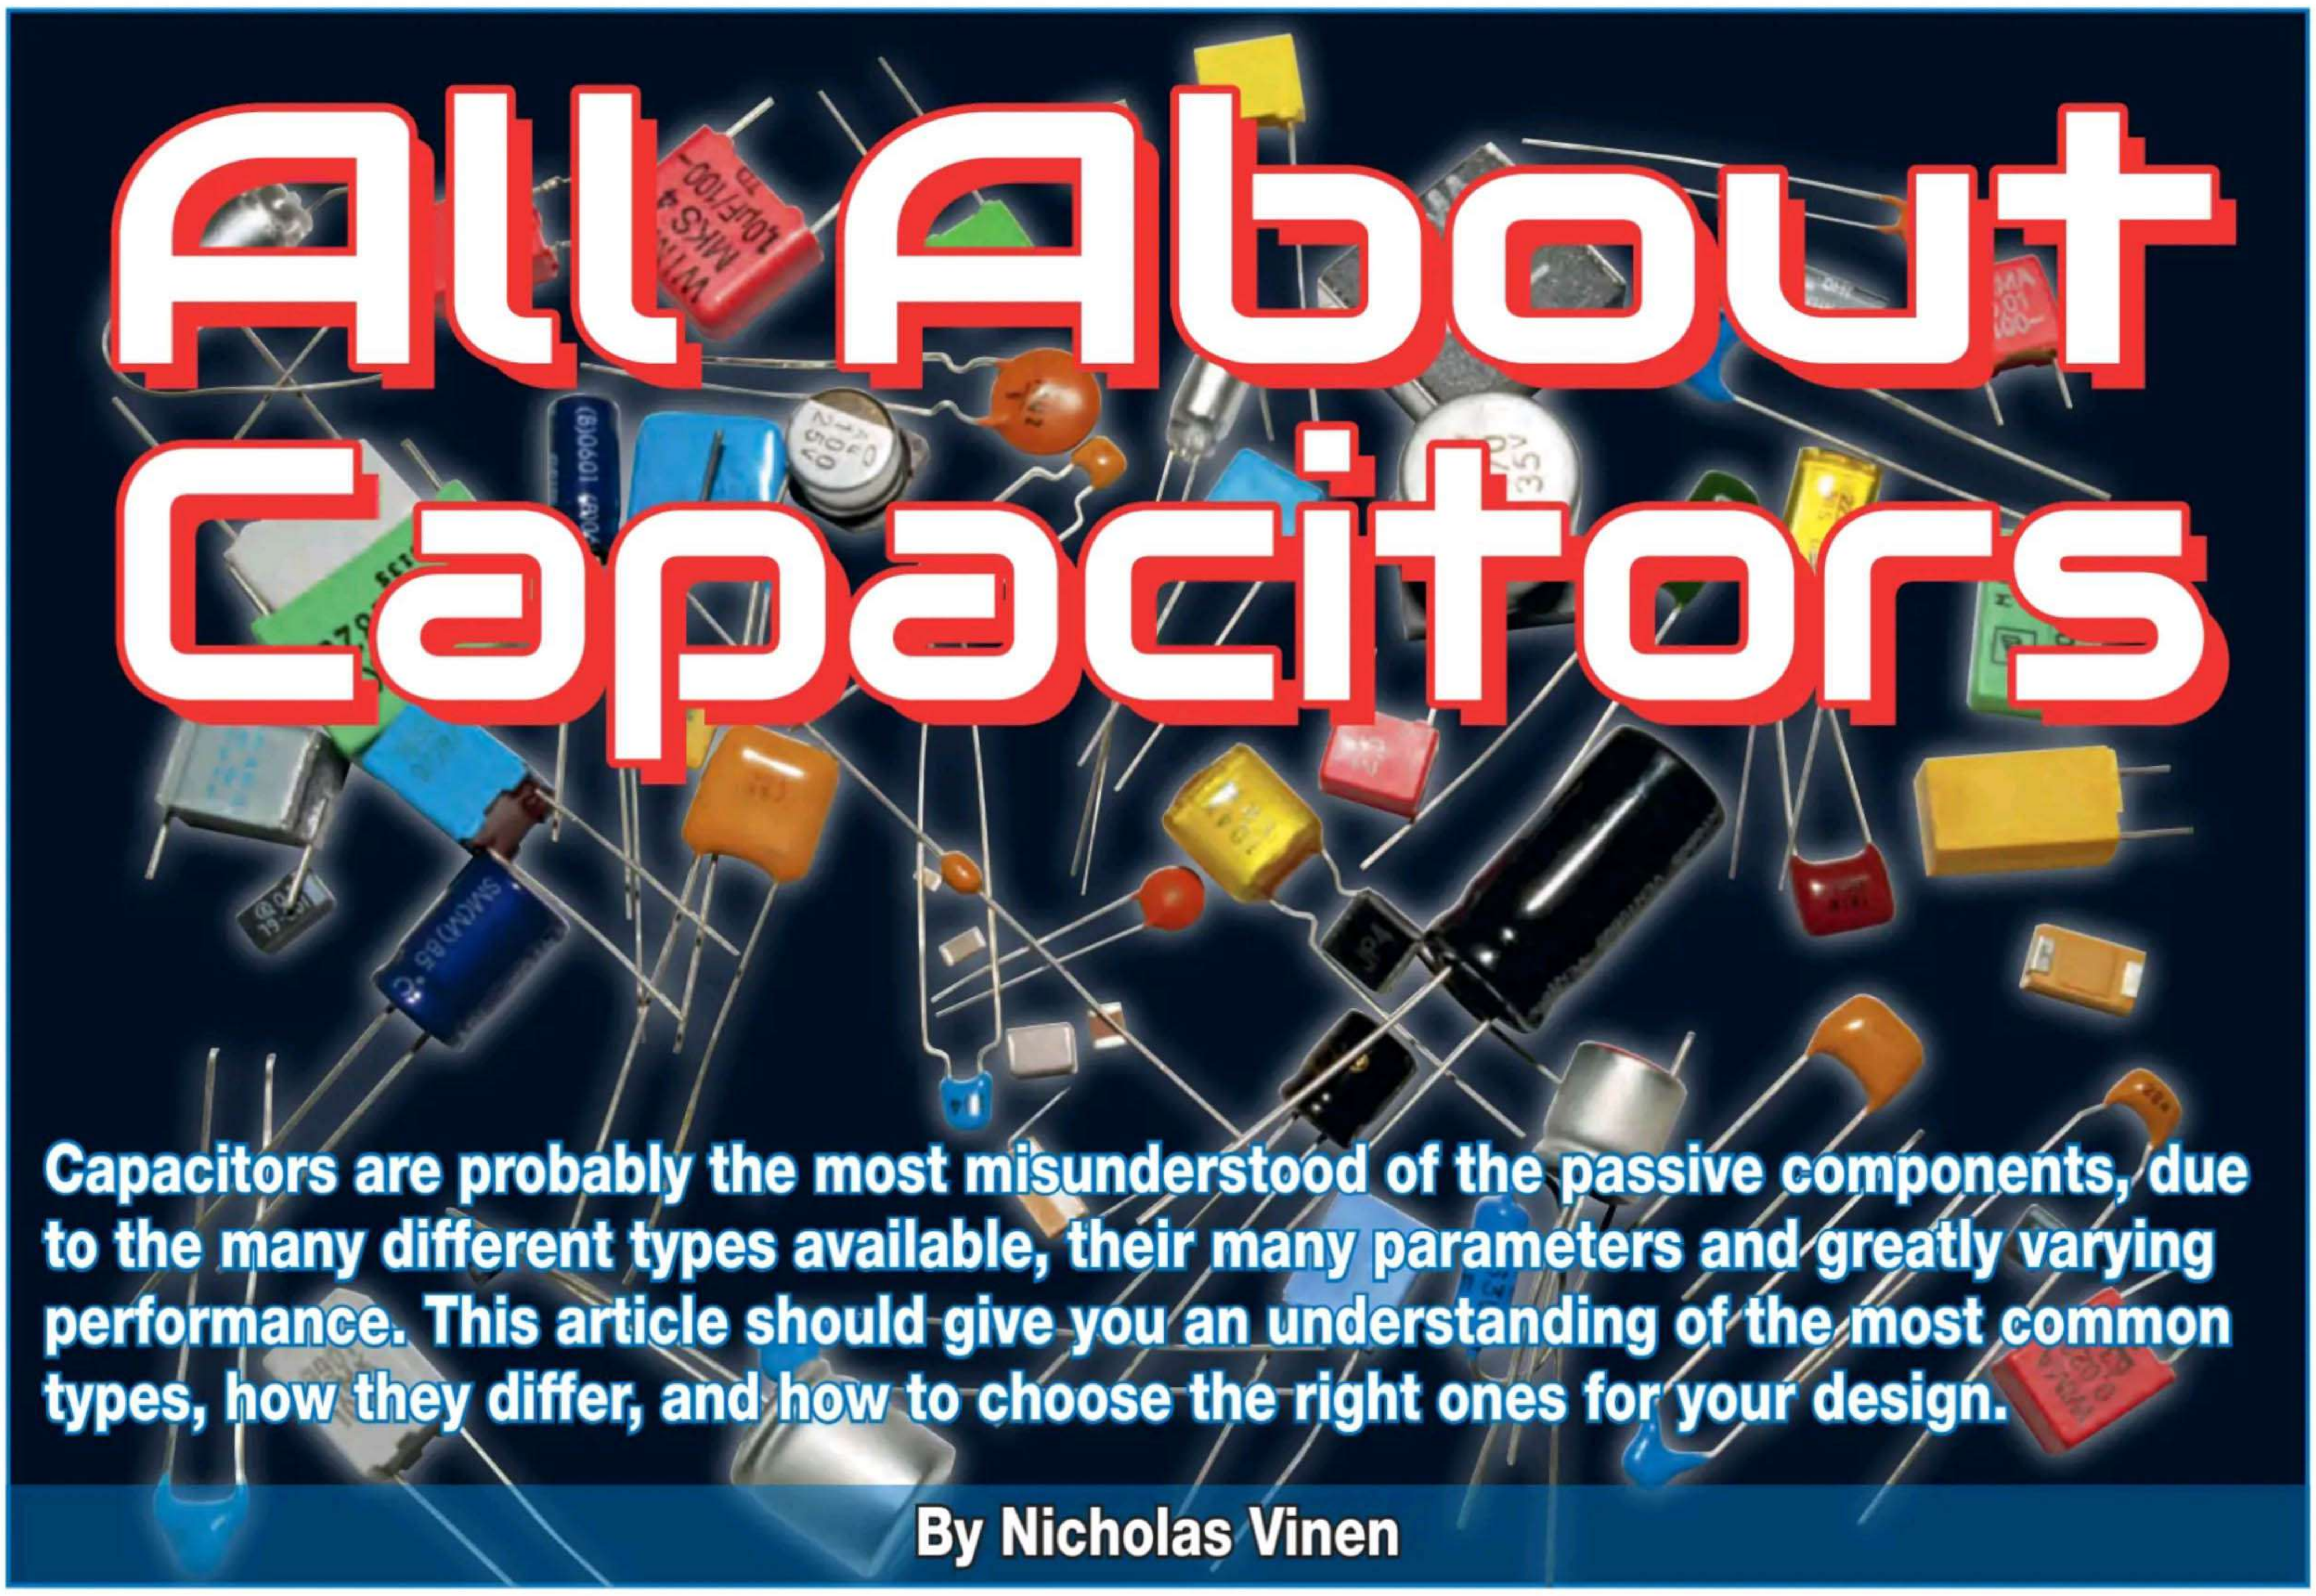

All About Capacitors

Capacitors are probably the most misunderstood of the passive components, due to the many different types available, their many parameters and greatly varying performance. This article should give you an understanding of the most common types, how they differ, and how to choose the right ones for your design.

By Nicholas Vinen

Capacitors come in all shapes and sizes. Some are much smaller than a grain of rice, while others are huge and used in banks to launch aircraft weighing many tonnes into the air!

Because there are so many different types, it can be very confusing trying to choose one. Even if you know what capacitance and voltage rating you need, there could be hundreds or even thousands of matching parts. Some of those might not work at all in your circuit, while others might work but not very well, and some will be very expensive. You need to narrow the choice down to just a handful and then pick one.

We have tried to break the following descriptions into digestible sections, despite their complexity. If you find yourself overwhelmed, give yourself time to digest what you have read so far, then read the rest later.

Capacitor dielectrics

Fundamentally, a capacitor is just two conductors (originally flat plates) separated by an insulator (the "dielectric"). But because the area of the plates required for any significant capacitance is quite large, modern capacitors are typically arranged as many layers of smaller conductors and insulators connected in parallel, allowing for a more compact package.

In some cases, the 'plates' are not even flat but instead are spiral coils, or 3D structures such as the etched surface of a metal foil or granular materials.

Etched or granular materials have a much higher capacitance per volume, as capacitance is proportional to surface area and inversely proportional to the distance between the plates.

This creates a tradeoff; thinner dielectrics give more capacitance, but have a lower breakdown voltage, so the maximum voltage applied to the capacitor must be kept lower. This is the main reason that a capacitor with a higher voltage rating, but the same capacitance, tends to be physically larger; its dielectric layer(s) need to be thicker.

The type of insulating (dielectric) material used has a strong effect on capacitor behaviour, and for this reason, capacitors are mostly categorised by the dielectric type. Different dielectric types have their own trade-offs in terms of capacitance, voltage ratings, linearity, current handling and more.

Some widely used dielectric materials for capacitors are:

- Ceramics (typically metal oxides)
- Metal oxide layers (in electrolytic capacitors)
- Plastic films
- Mica
- Paper
- A Helmholtz plane of solvent molecules (as in 'double layer' super/ultracapacitors)

The most common types of capacitors in use today are ceramic and electrolytic, followed by plastic film types. These three types of capacitors have important sub-categories which strongly affect their behaviour.

One property of all dielectric materials is the dielectric constant ("K"). The larger this number, the higher the capacitance for a similarly constructed device. K can vary with temperature, voltage, age and other properties. While high K values make for greater capacitances in a small volume, there are significant penalties in other areas, as we describe below.

All About Capacitors

Ceramic capacitors

If you look at the PCB of just about any modern electronic device, you will find it covered in ceramic capacitors. They are cheap, reliable, perform very well and are available in a wide range of capacitances and voltage ratings.

Because modern ceramic capacitors are fabricated in bulk, they can have anywhere from one to many thousands of layers. This gives them a wide capacitance range, from fractions of a picofarad up to hundreds of microfarads, in a small package – see Figs.1-3.

Ceramic capacitors are typically robust and long-lasting, and are not polarised (they can handle negative or positive voltages).

Ceramic capacitors are available with voltage ratings from just a few volts up to several kilovolts. Ceramic capacitors with voltage ratings above 500V tend to use different types of ceramic to those below 500V, and have slightly different properties.

The most common ceramics used are based on titanium dioxide (TiO_2) or barium titanate ($BaTiO_3$) with additives to tweak their properties.

As there are so many different possible combinations, they are arranged in various categories based on their performance. The categories are based on the initial tolerance of the capacitor (ie, the variation of real samples from the rated value), how the capacitance changes with temperature (the temperature coefficient) and how it changes with applied voltage (the voltage coefficient).

The most common type codes are NP0 or C0G (different names for the same category), JB, SL0, X5R, X5S, X6S, X7R, X7S, X8L, Y5V and Z5U.

To take three examples, NP0/C0G types have very close tolerances and no or minimal capacitance variation with temperature or voltage. They also have a low dielectric constant, so they are relatively large for a given capacitance value and voltage rating. As a result, they are also quite expensive.

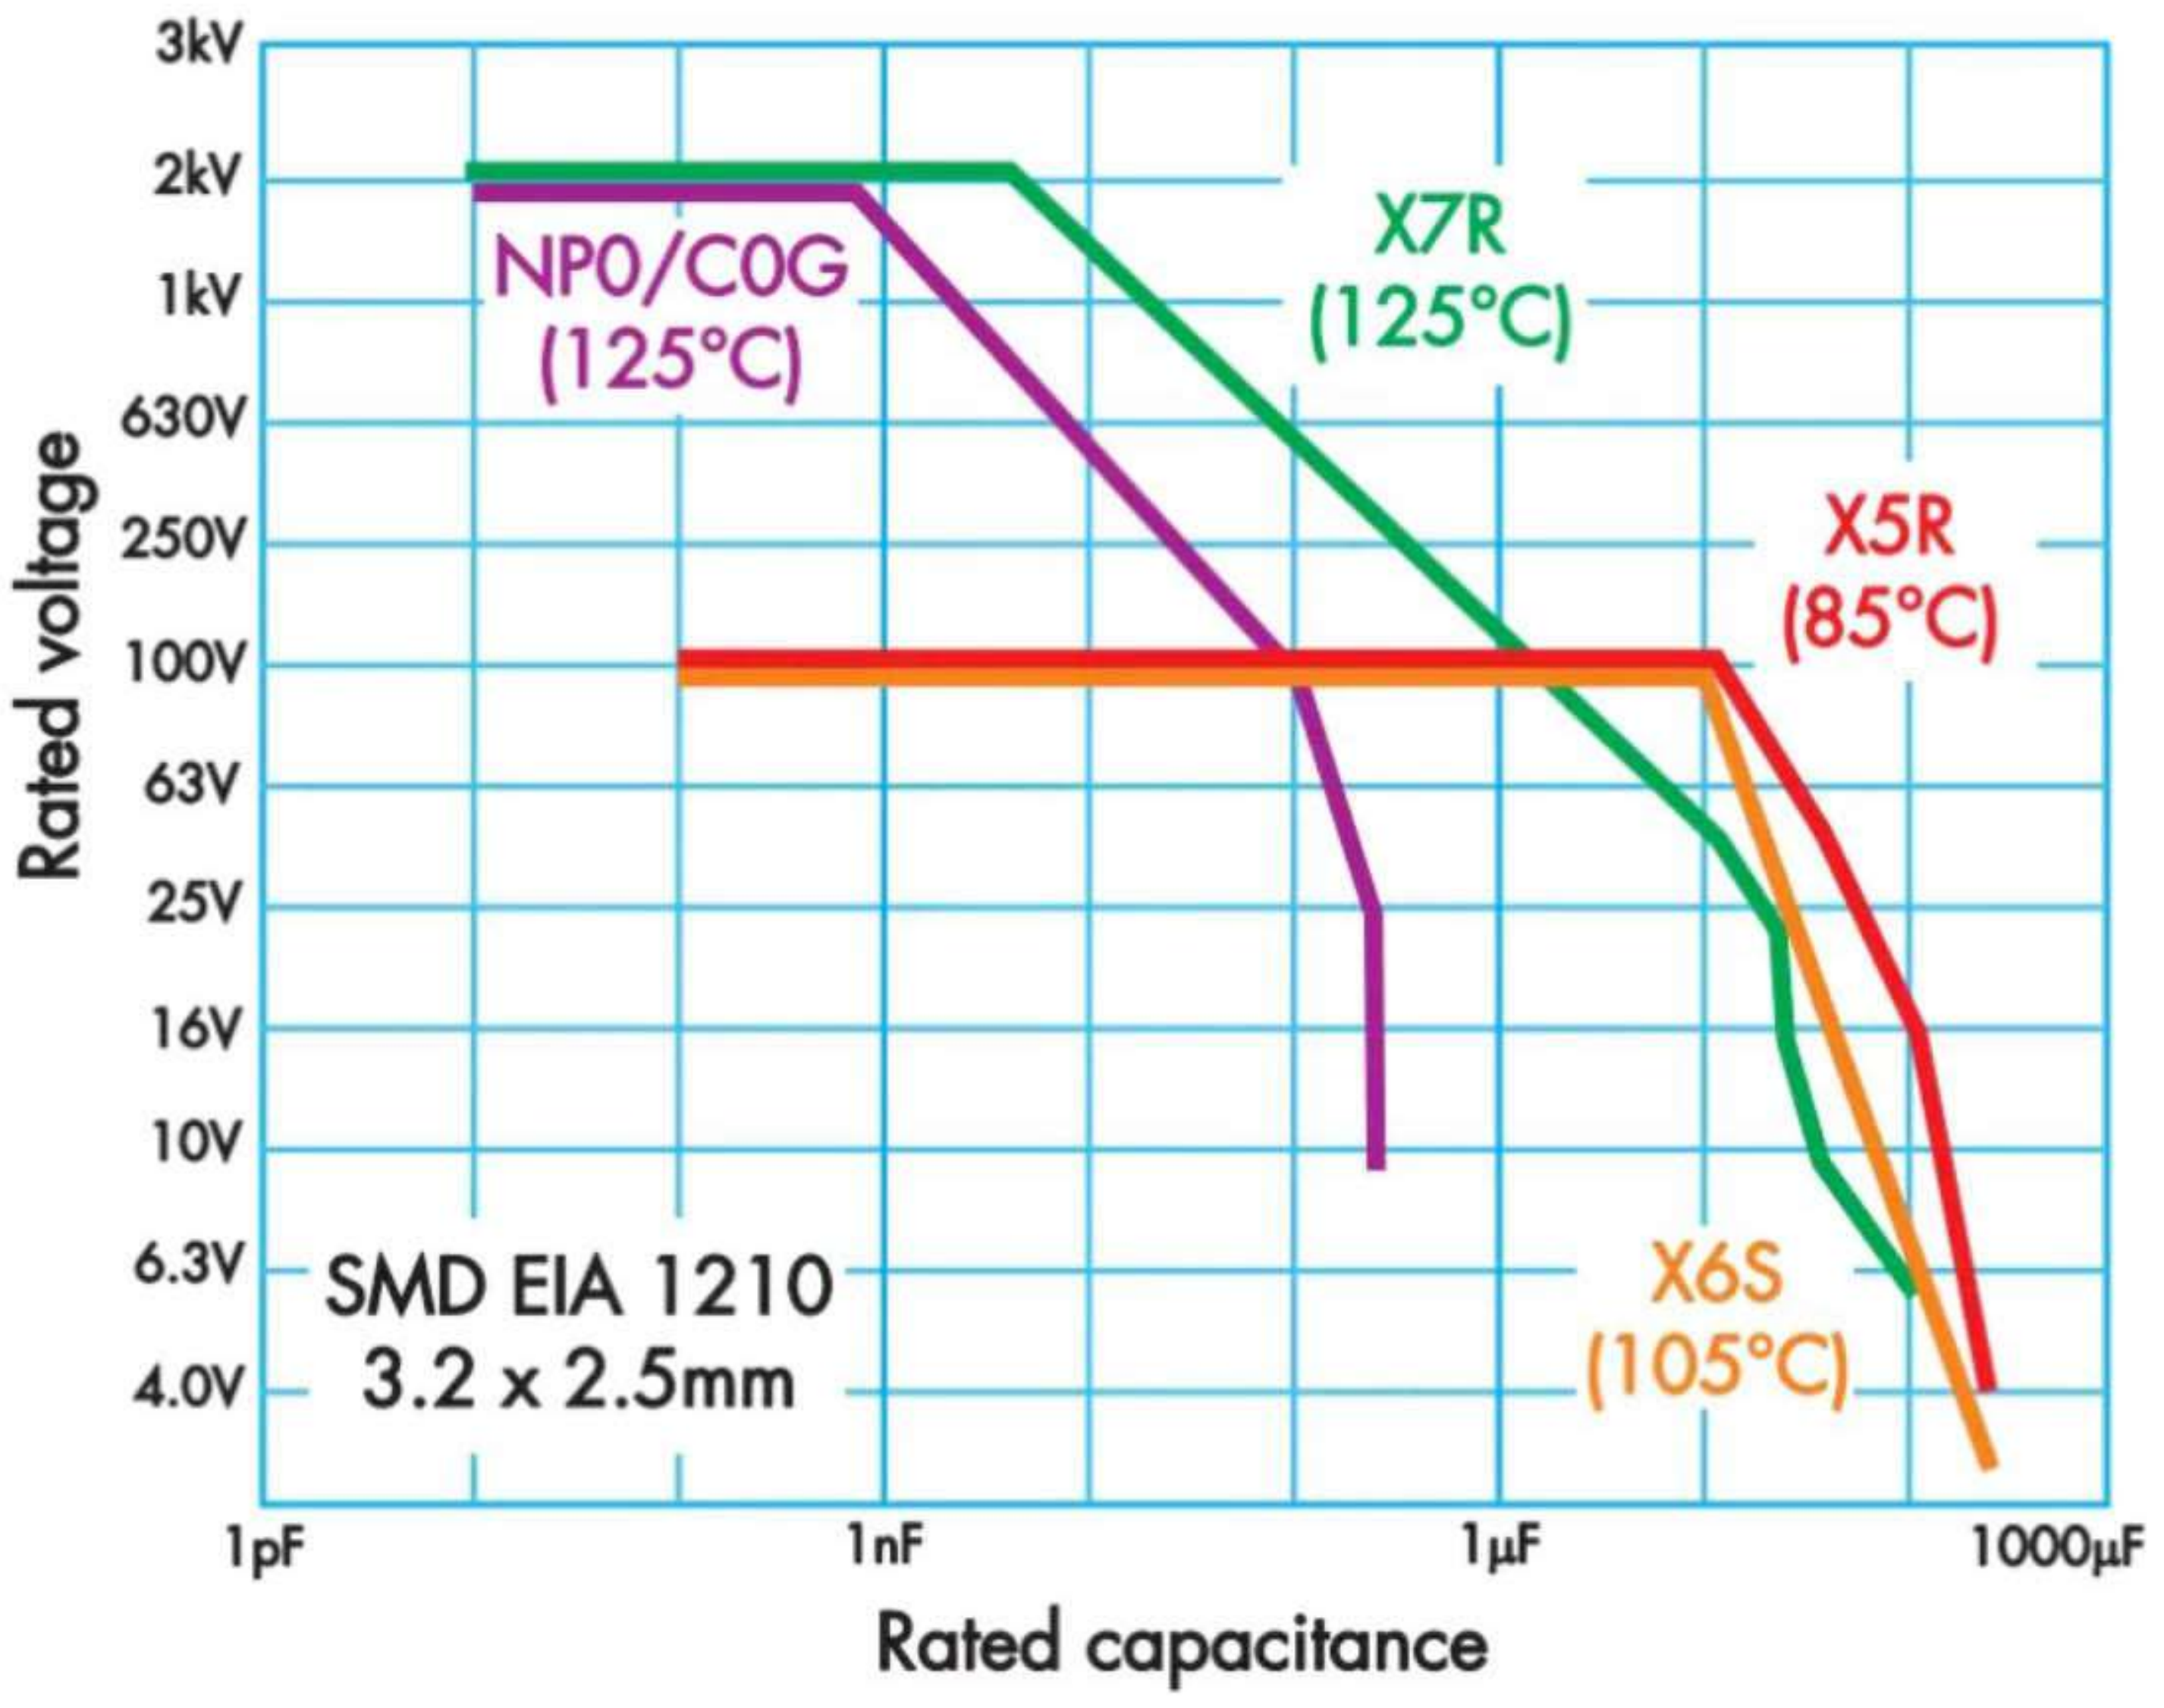

Fig.1: the range of capacitances and voltages available in 3.2 x 2.5mm SMD ceramic capacitors today. Both larger and smaller sizes are available, extending the range of values down to 0.1pF (1.6 x 0.8mm) and up to 470µF (4.5 x 3.2mm). Note how some types of ceramic dielectric are available to higher working voltages, and others to a higher maximum capacitance. *(original source: Wikipedia)*

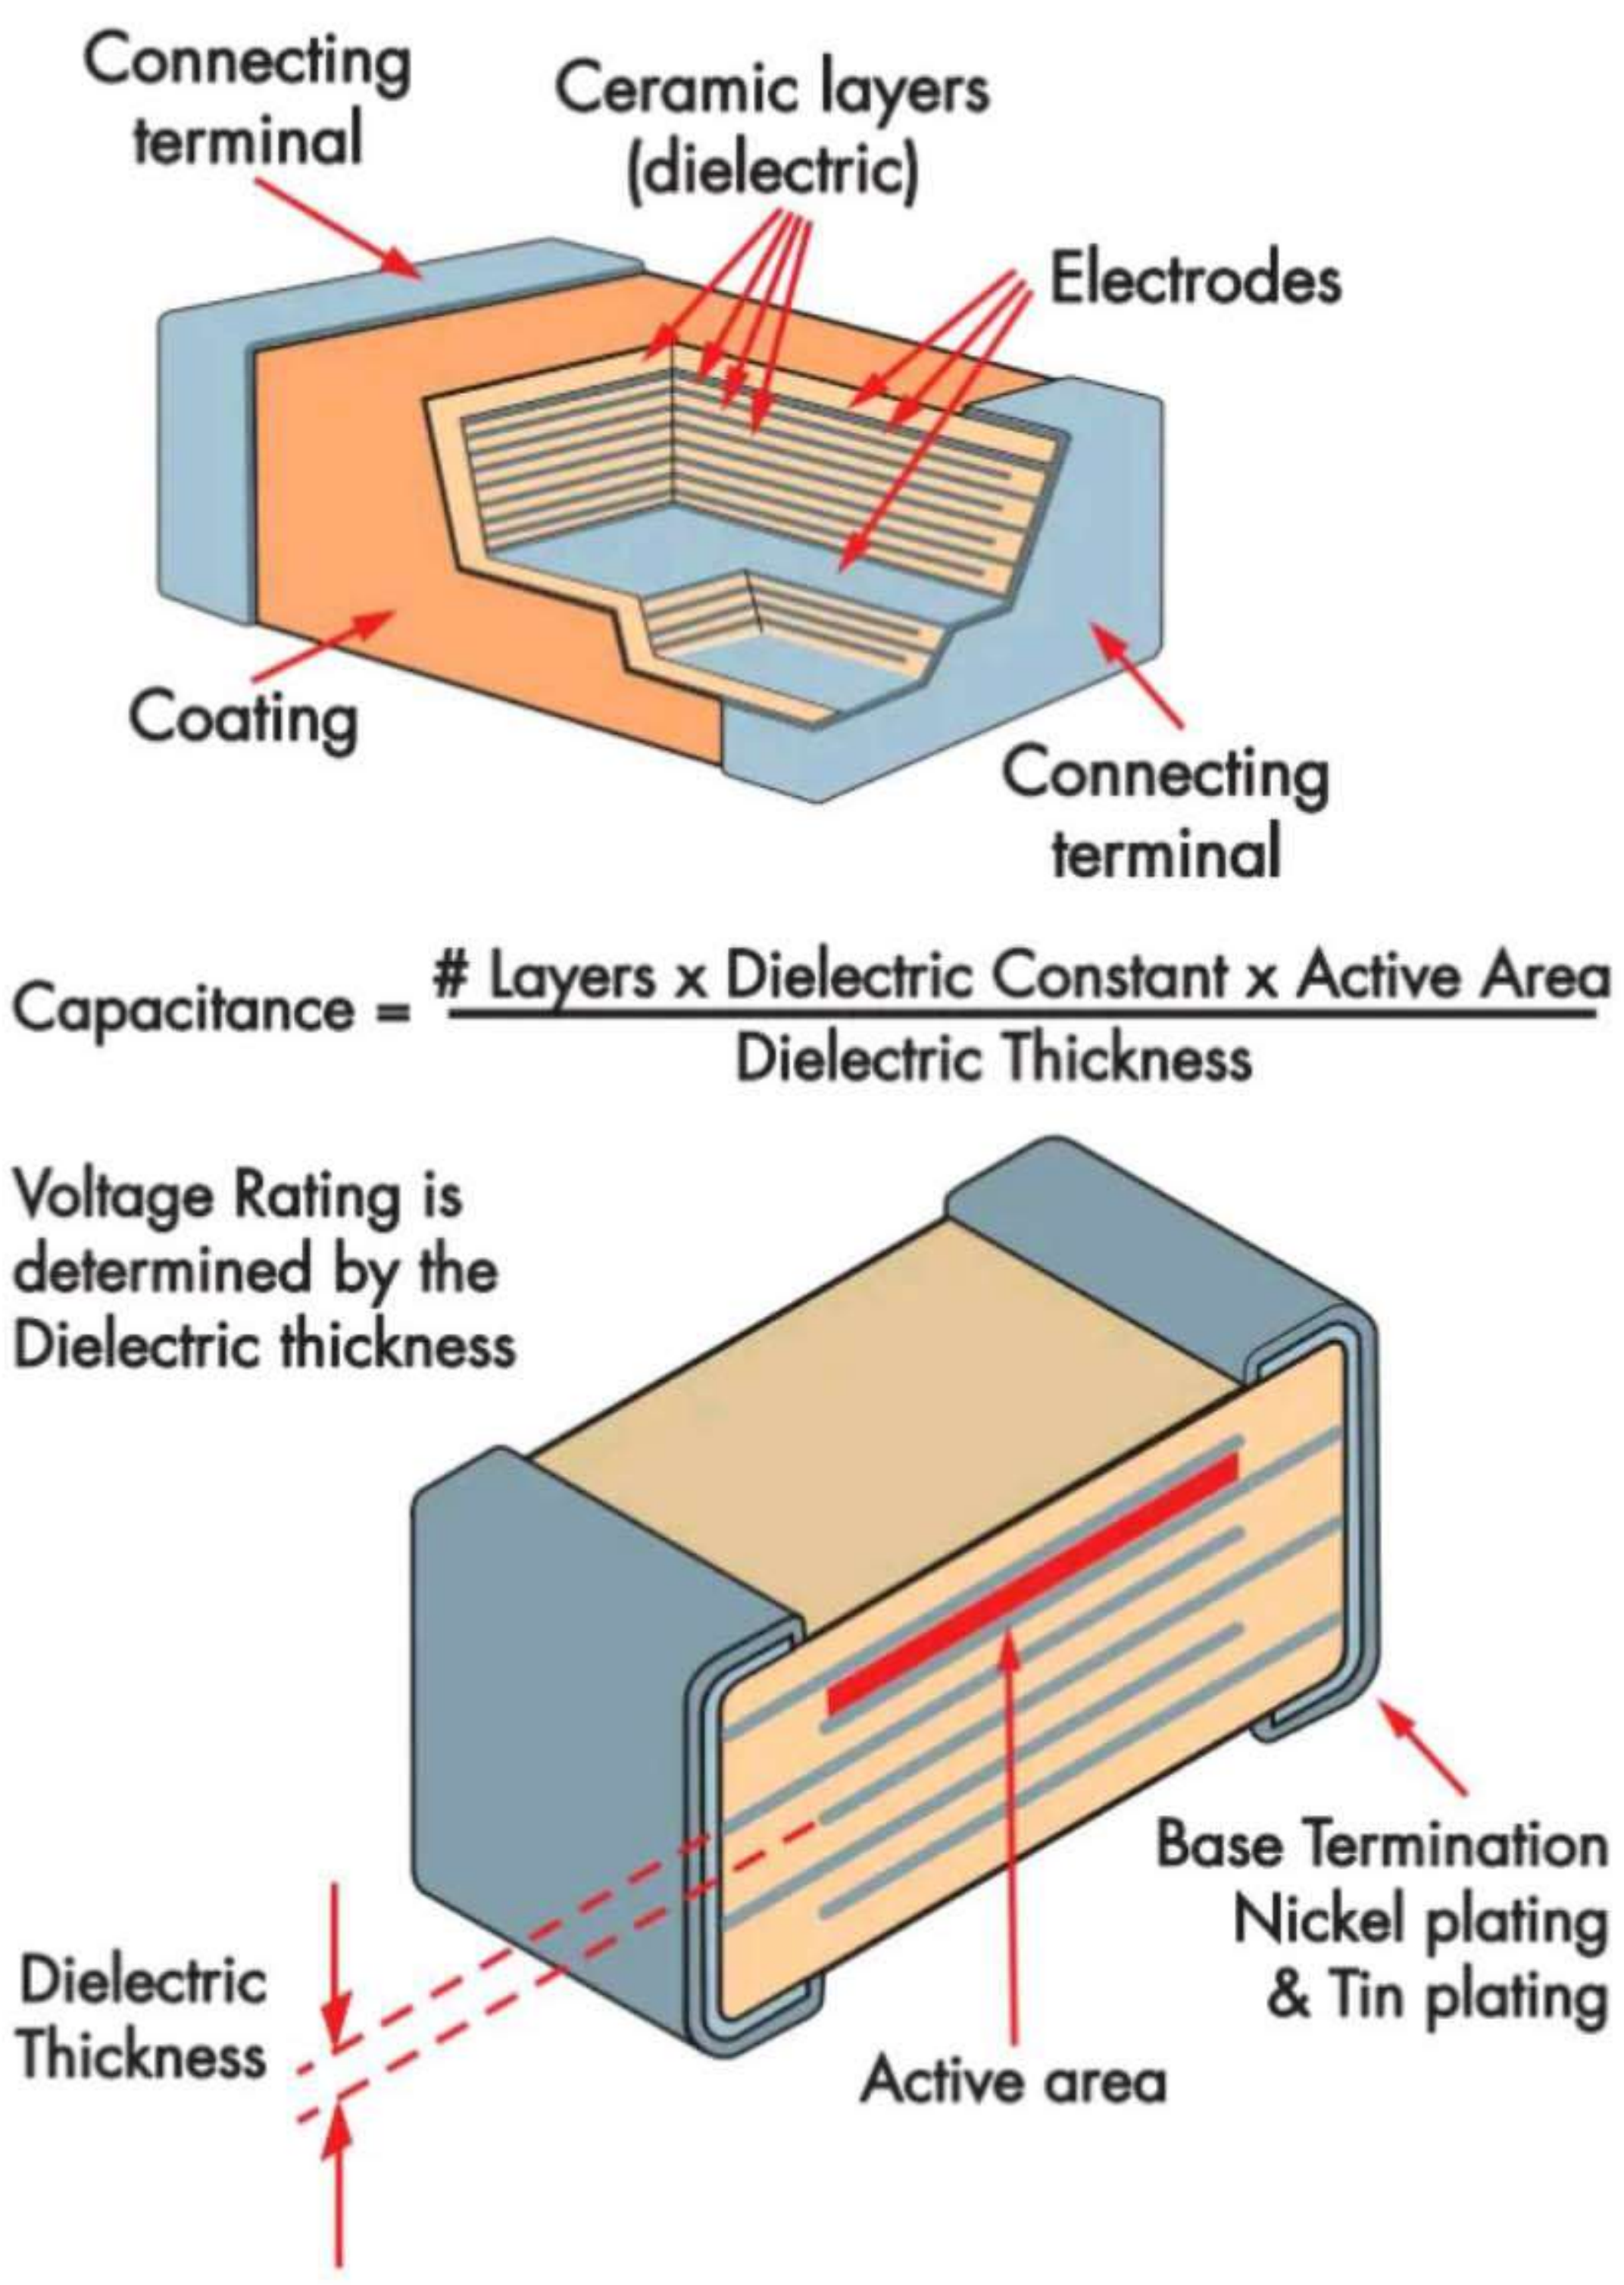
Connecting terminal
Ceramic layers (dielectric)
Electrodes
Coating
Connecting terminal
Capacitance = # Layers x Dielectric Constant x Active Area / Dielectric Thickness
Voltage Rating is determined by the Dielectric thickness
Base Termination Nickel plating & Tin plating
Dielectric Thickness
Active area

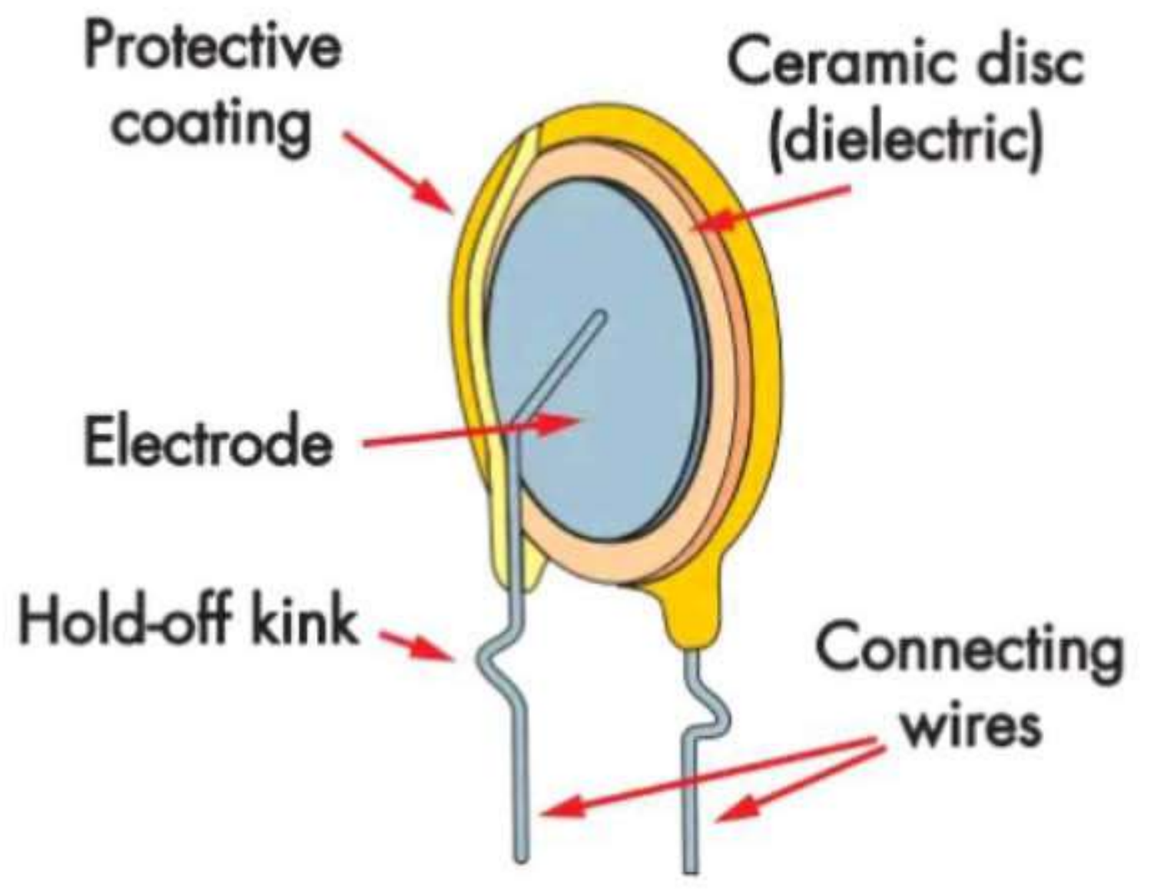
Protective coating
Ceramic disc (dielectric)
Electrode
Hold-off kink
Connecting wires

Fig.2: the structure of typical SMD and through-hole ceramic capacitors. SMD 'chip' ceramics are made of many layers; through-hole disc capacitors may have a single layer construction, as shown here, or increasingly these days, a similar internal structure to a multi-layer SMD capacitor. Multi-layer through-hole capacitors are usually encapsulated in epoxy, while the single-layer disc types can be encapsulated in ceramic. *(original source: Johanson Dielectrics)*

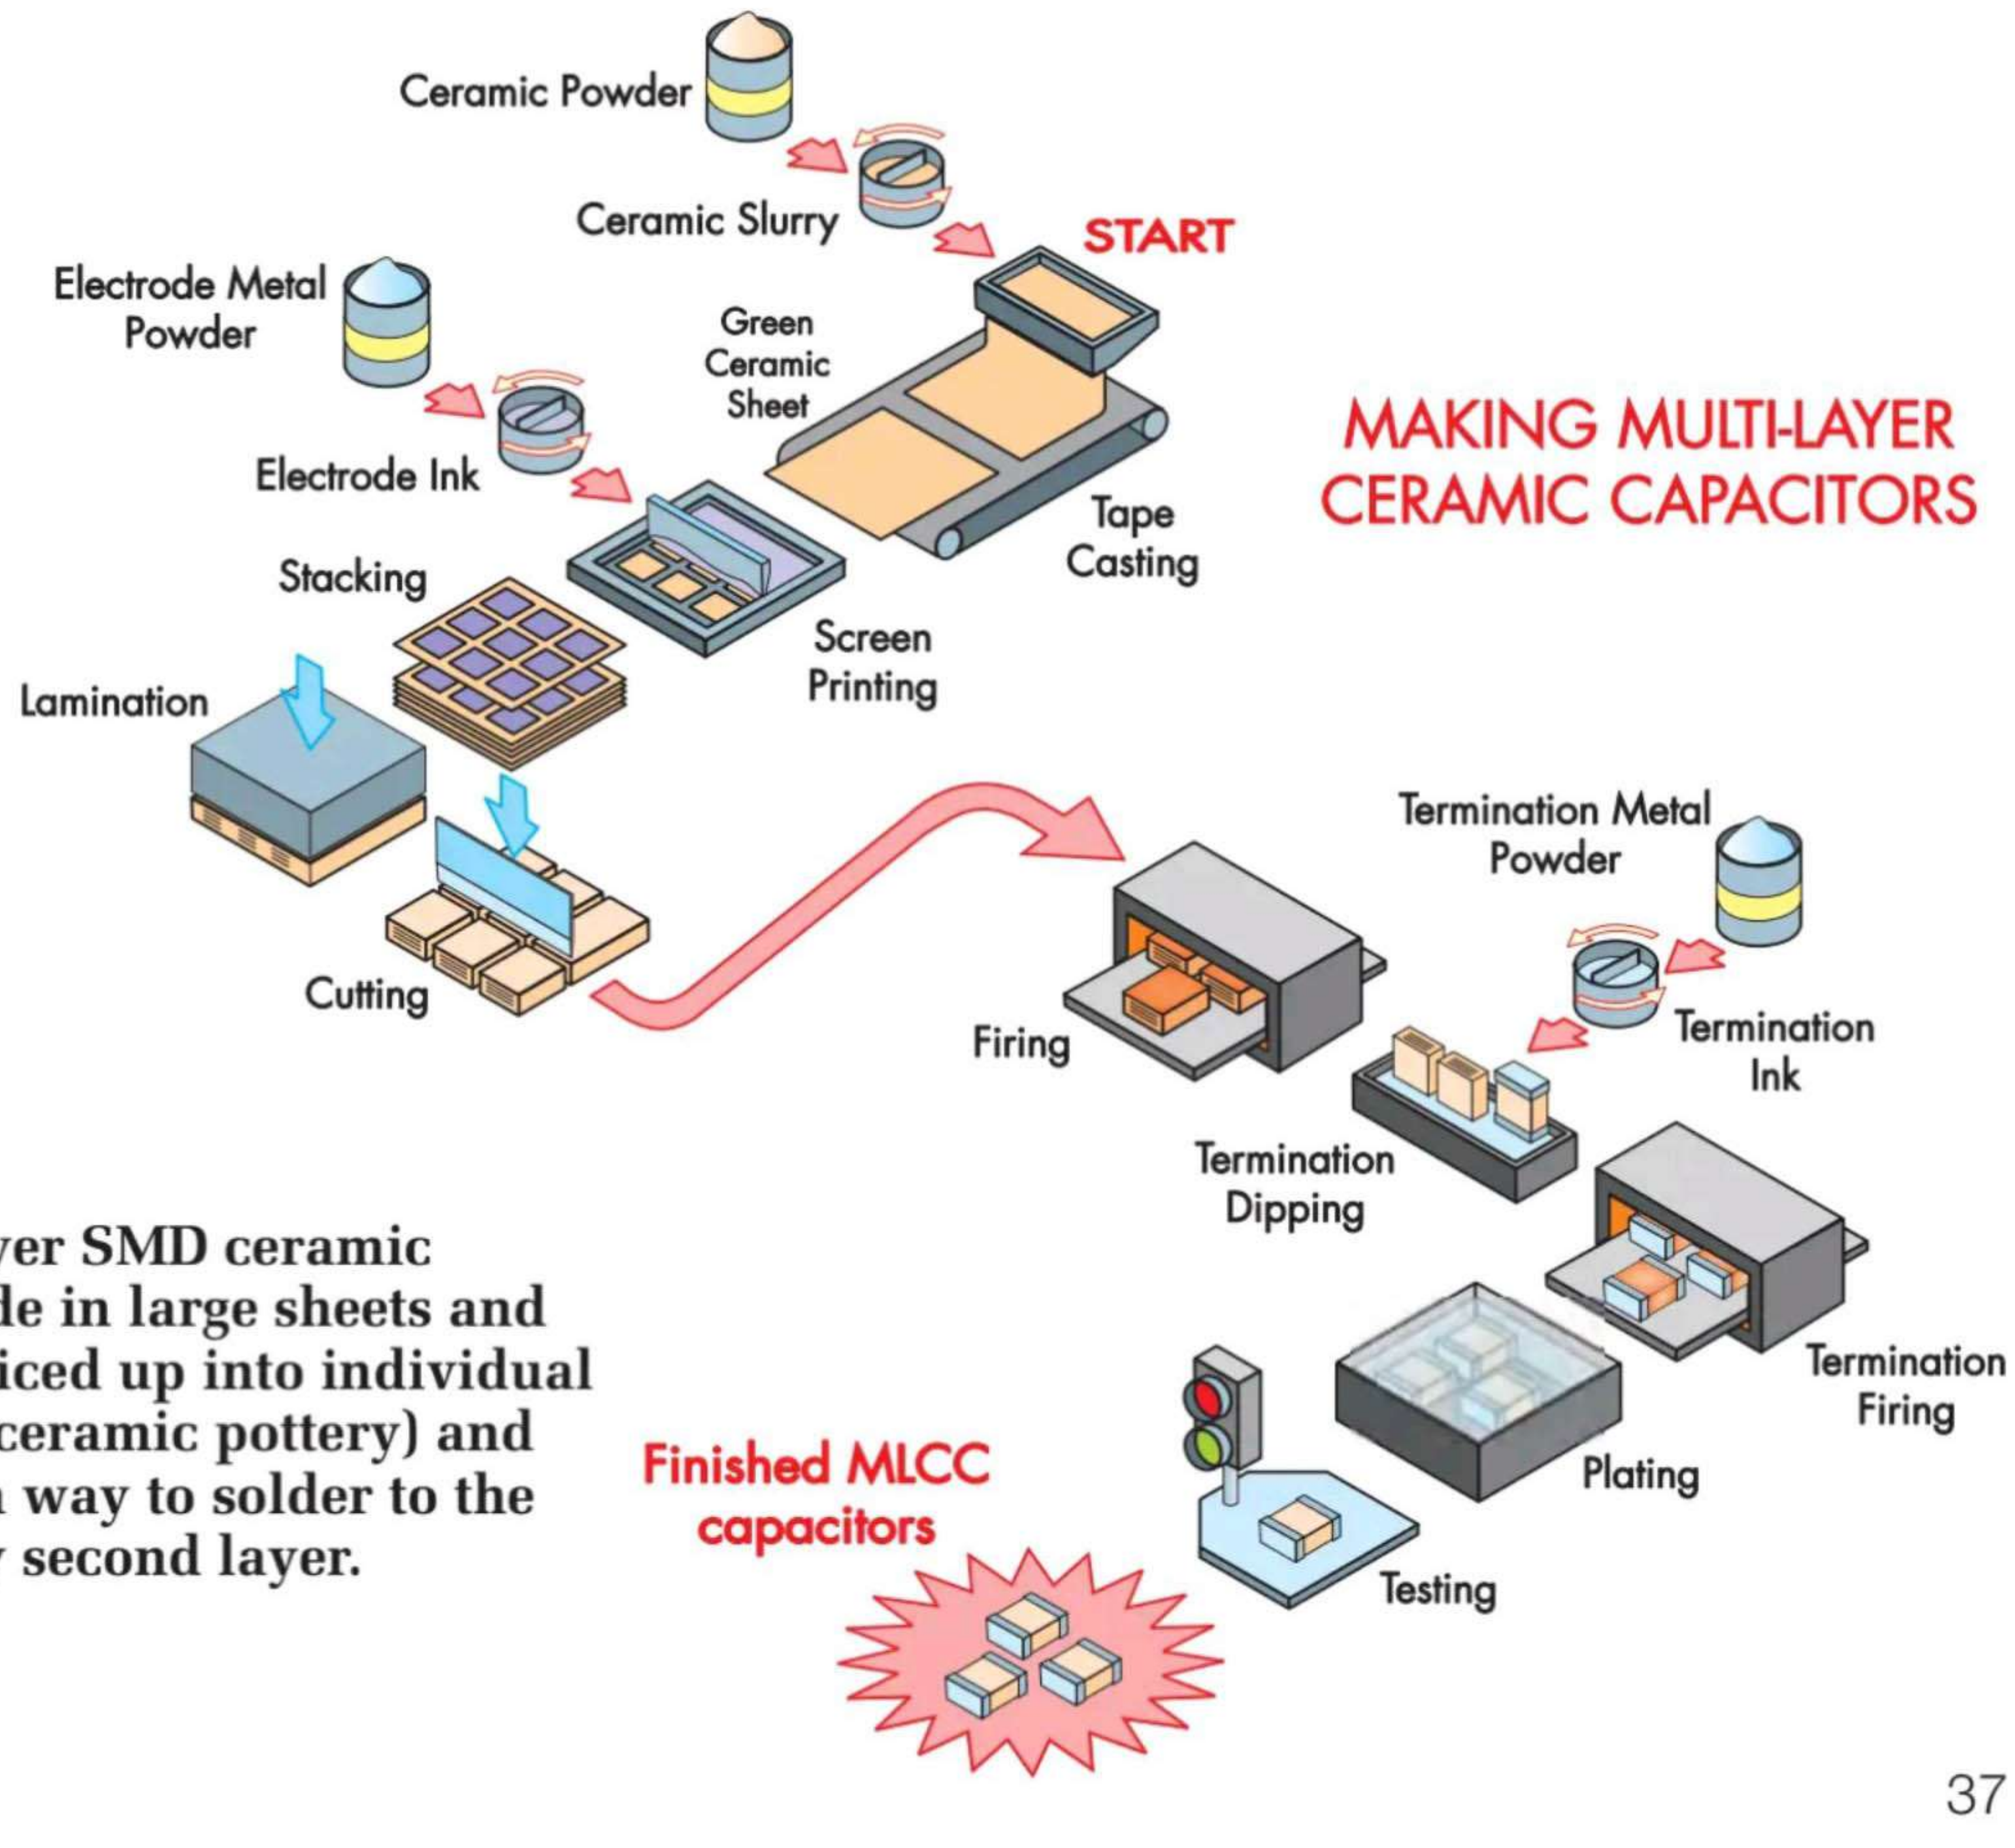
Ceramic Powder
Ceramic Slurry
START
Electrode Metal Powder
Green Ceramic Sheet
MAKING MULTI-LAYER CERAMIC CAPACITORS
Electrode Ink
Tape Casting
Stacking
Screen Printing
Lamination
Cutting
Termination Metal Powder
Firing
Termination Ink
Termination Dipping
Termination Firing
Finished MLCC capacitors
Plating
Testing

Fig.3: the manufacturing process for multi-layer SMD ceramic capacitors. To keep the cost low, they are made in large sheets and after lamination is complete, the sheets are sliced up into individual capacitors. Those are then fired (similarly to ceramic pottery) and the end terminals are added, which provide a way to solder to the capacitor while also electrically joining every second layer. *(original source: Johanson Dielectrics)*

103
103
104Z
102
104
47
220

A set of ceramic capacitors ranging from 47pF to 2.2µF.

Table 1: Class 2 capacitor codes

First letter lower temperature	Middle number upper temperature	Last letter change in capacitance over given temperature range
X = -55°C	4 = +65°C	P = ±10%
Y = -30°C	5 = +85°C	R = ±15%
Z = +10°C	6 = +105°C	L = ±15%, +15 / -40% above 125°C
	7 = +125°C	S = ±22%
	8 = +150°C	T = +22 / -33%
	9 = +200°C	U = +22 / -56%
		V = +22 / -82%

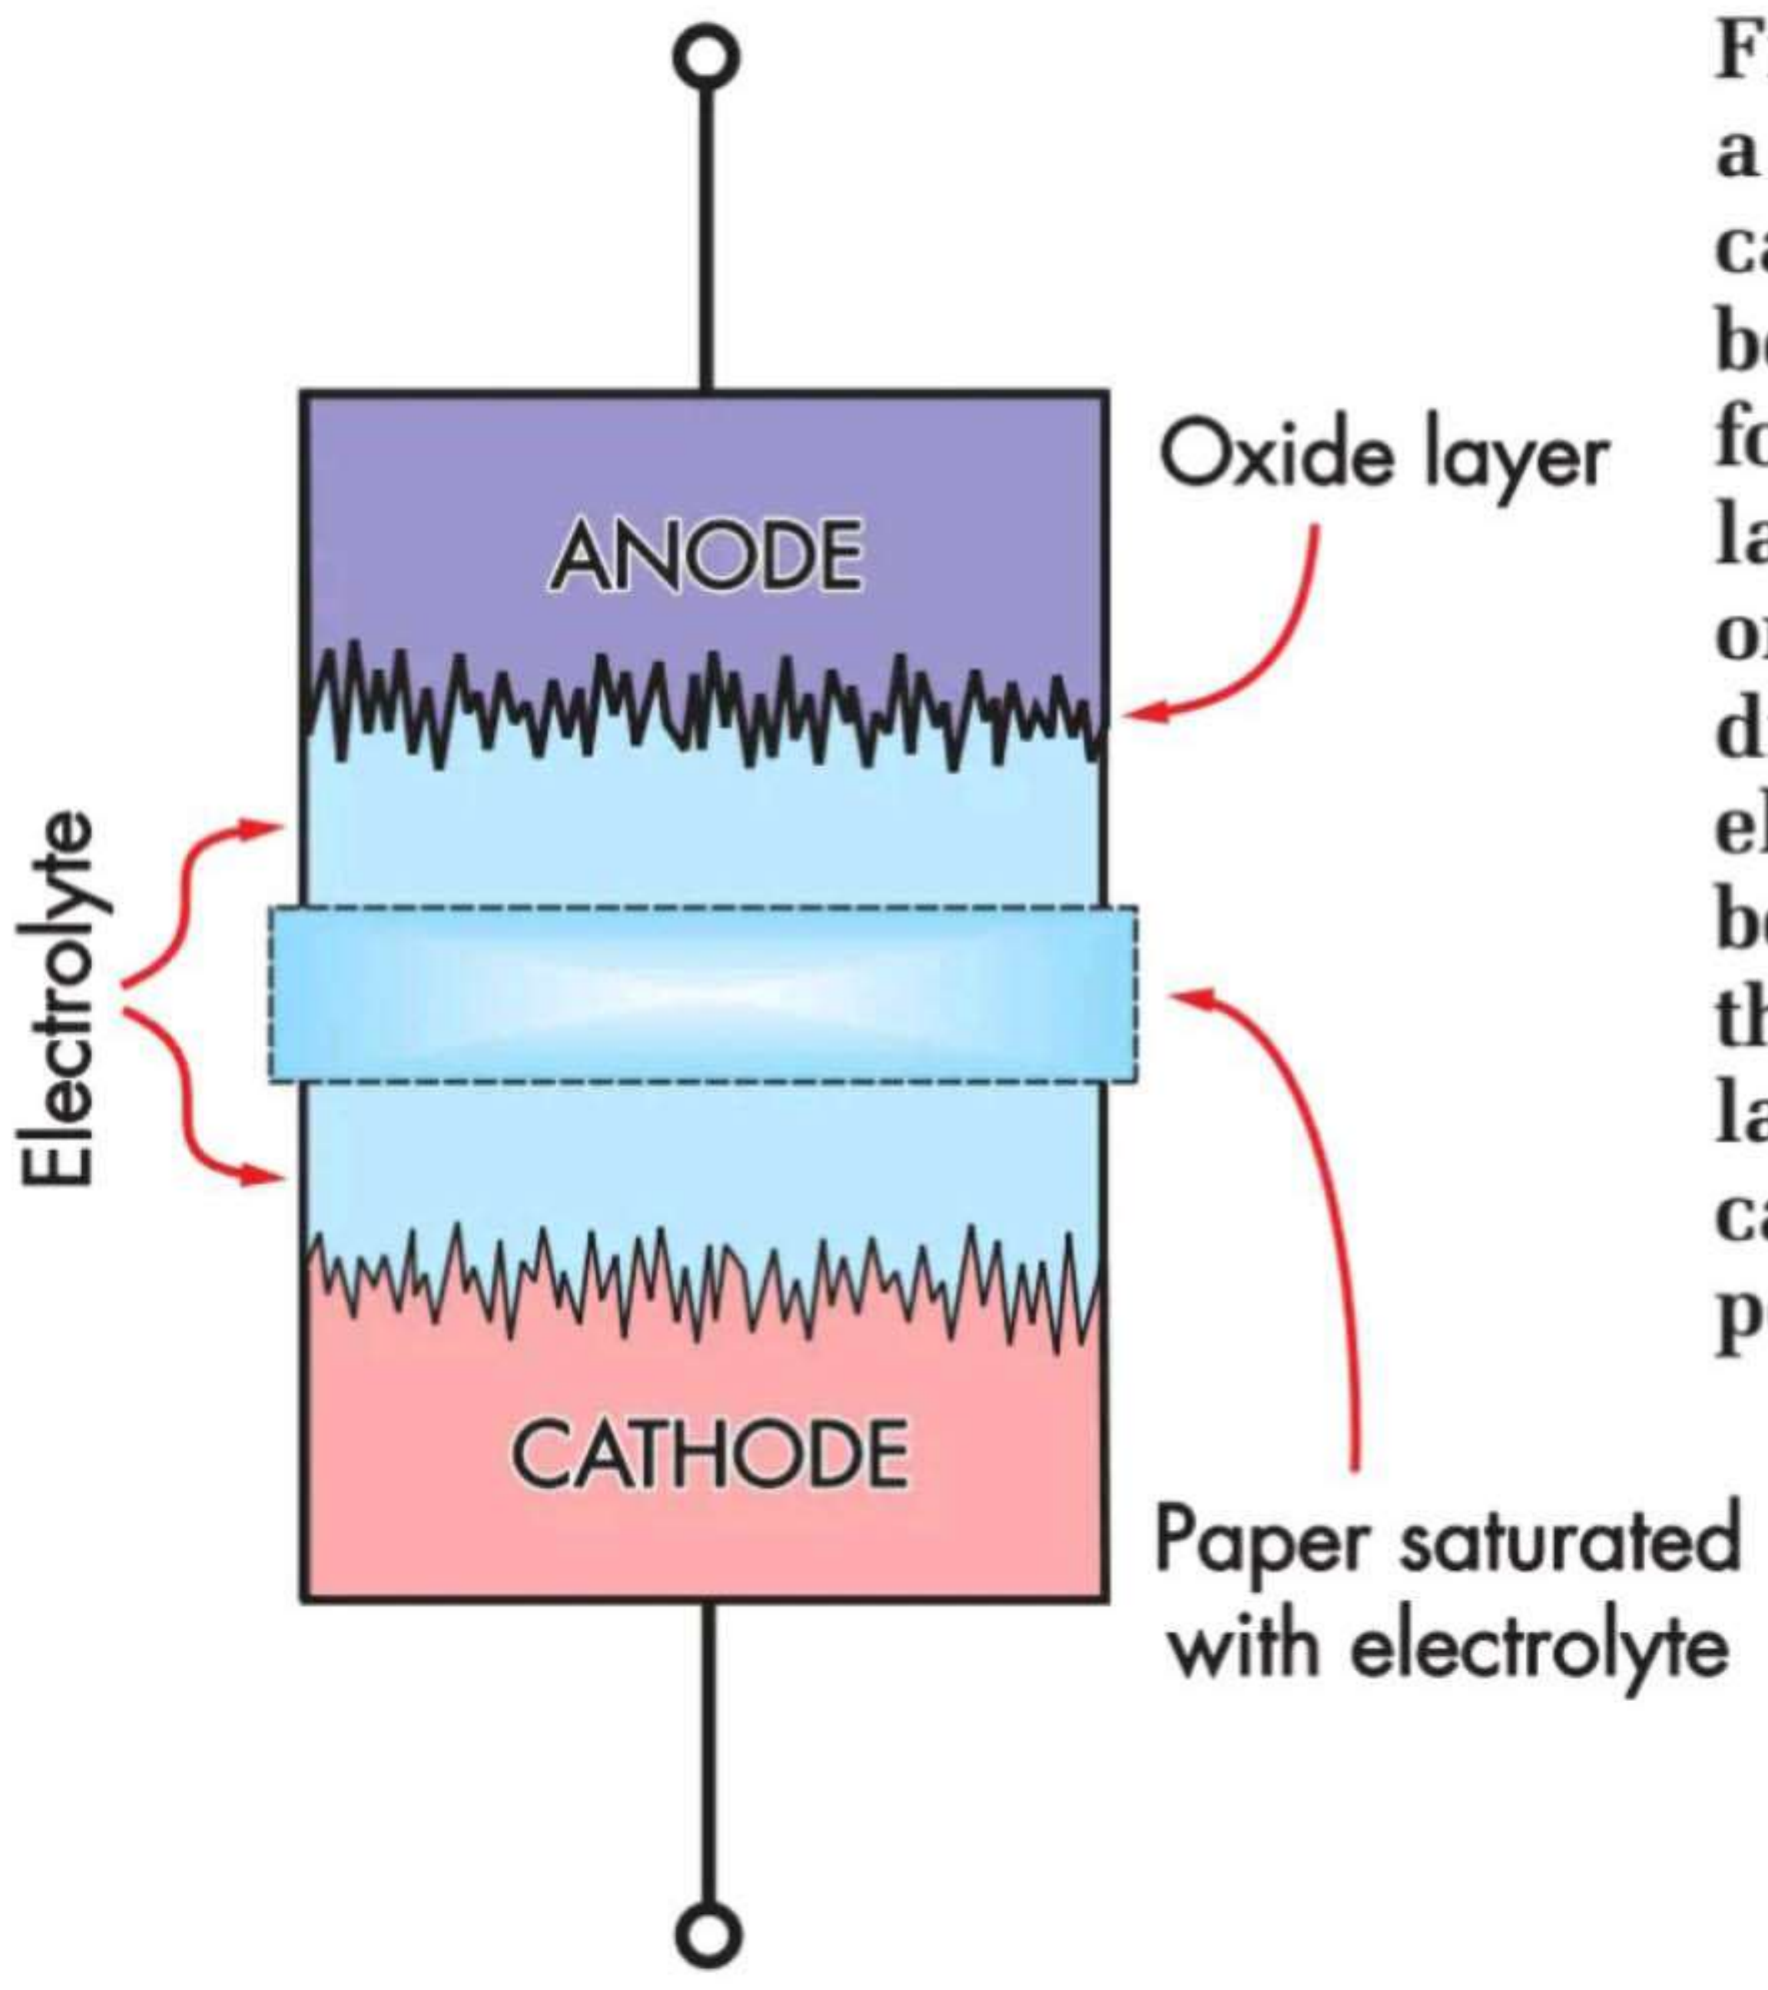
ANODE
Oxide layer
Electrolyte
CATHODE
Paper saturated with electrolyte

Fig.4: a cross-section of one layer of a standard aluminium electrolytic capacitor. The anode and cathode are both made from etched aluminium foil, for a large surface area. A thin layer of aluminium oxide is formed on the anode, which acts as the dielectric layer. The conductive electrolyte allows electrons to flow between the cathode right up against that oxide layer, so only the oxide layer separates the two halves of the capacitor, maximising capacitance per area. *(original source: Wikipedia)*

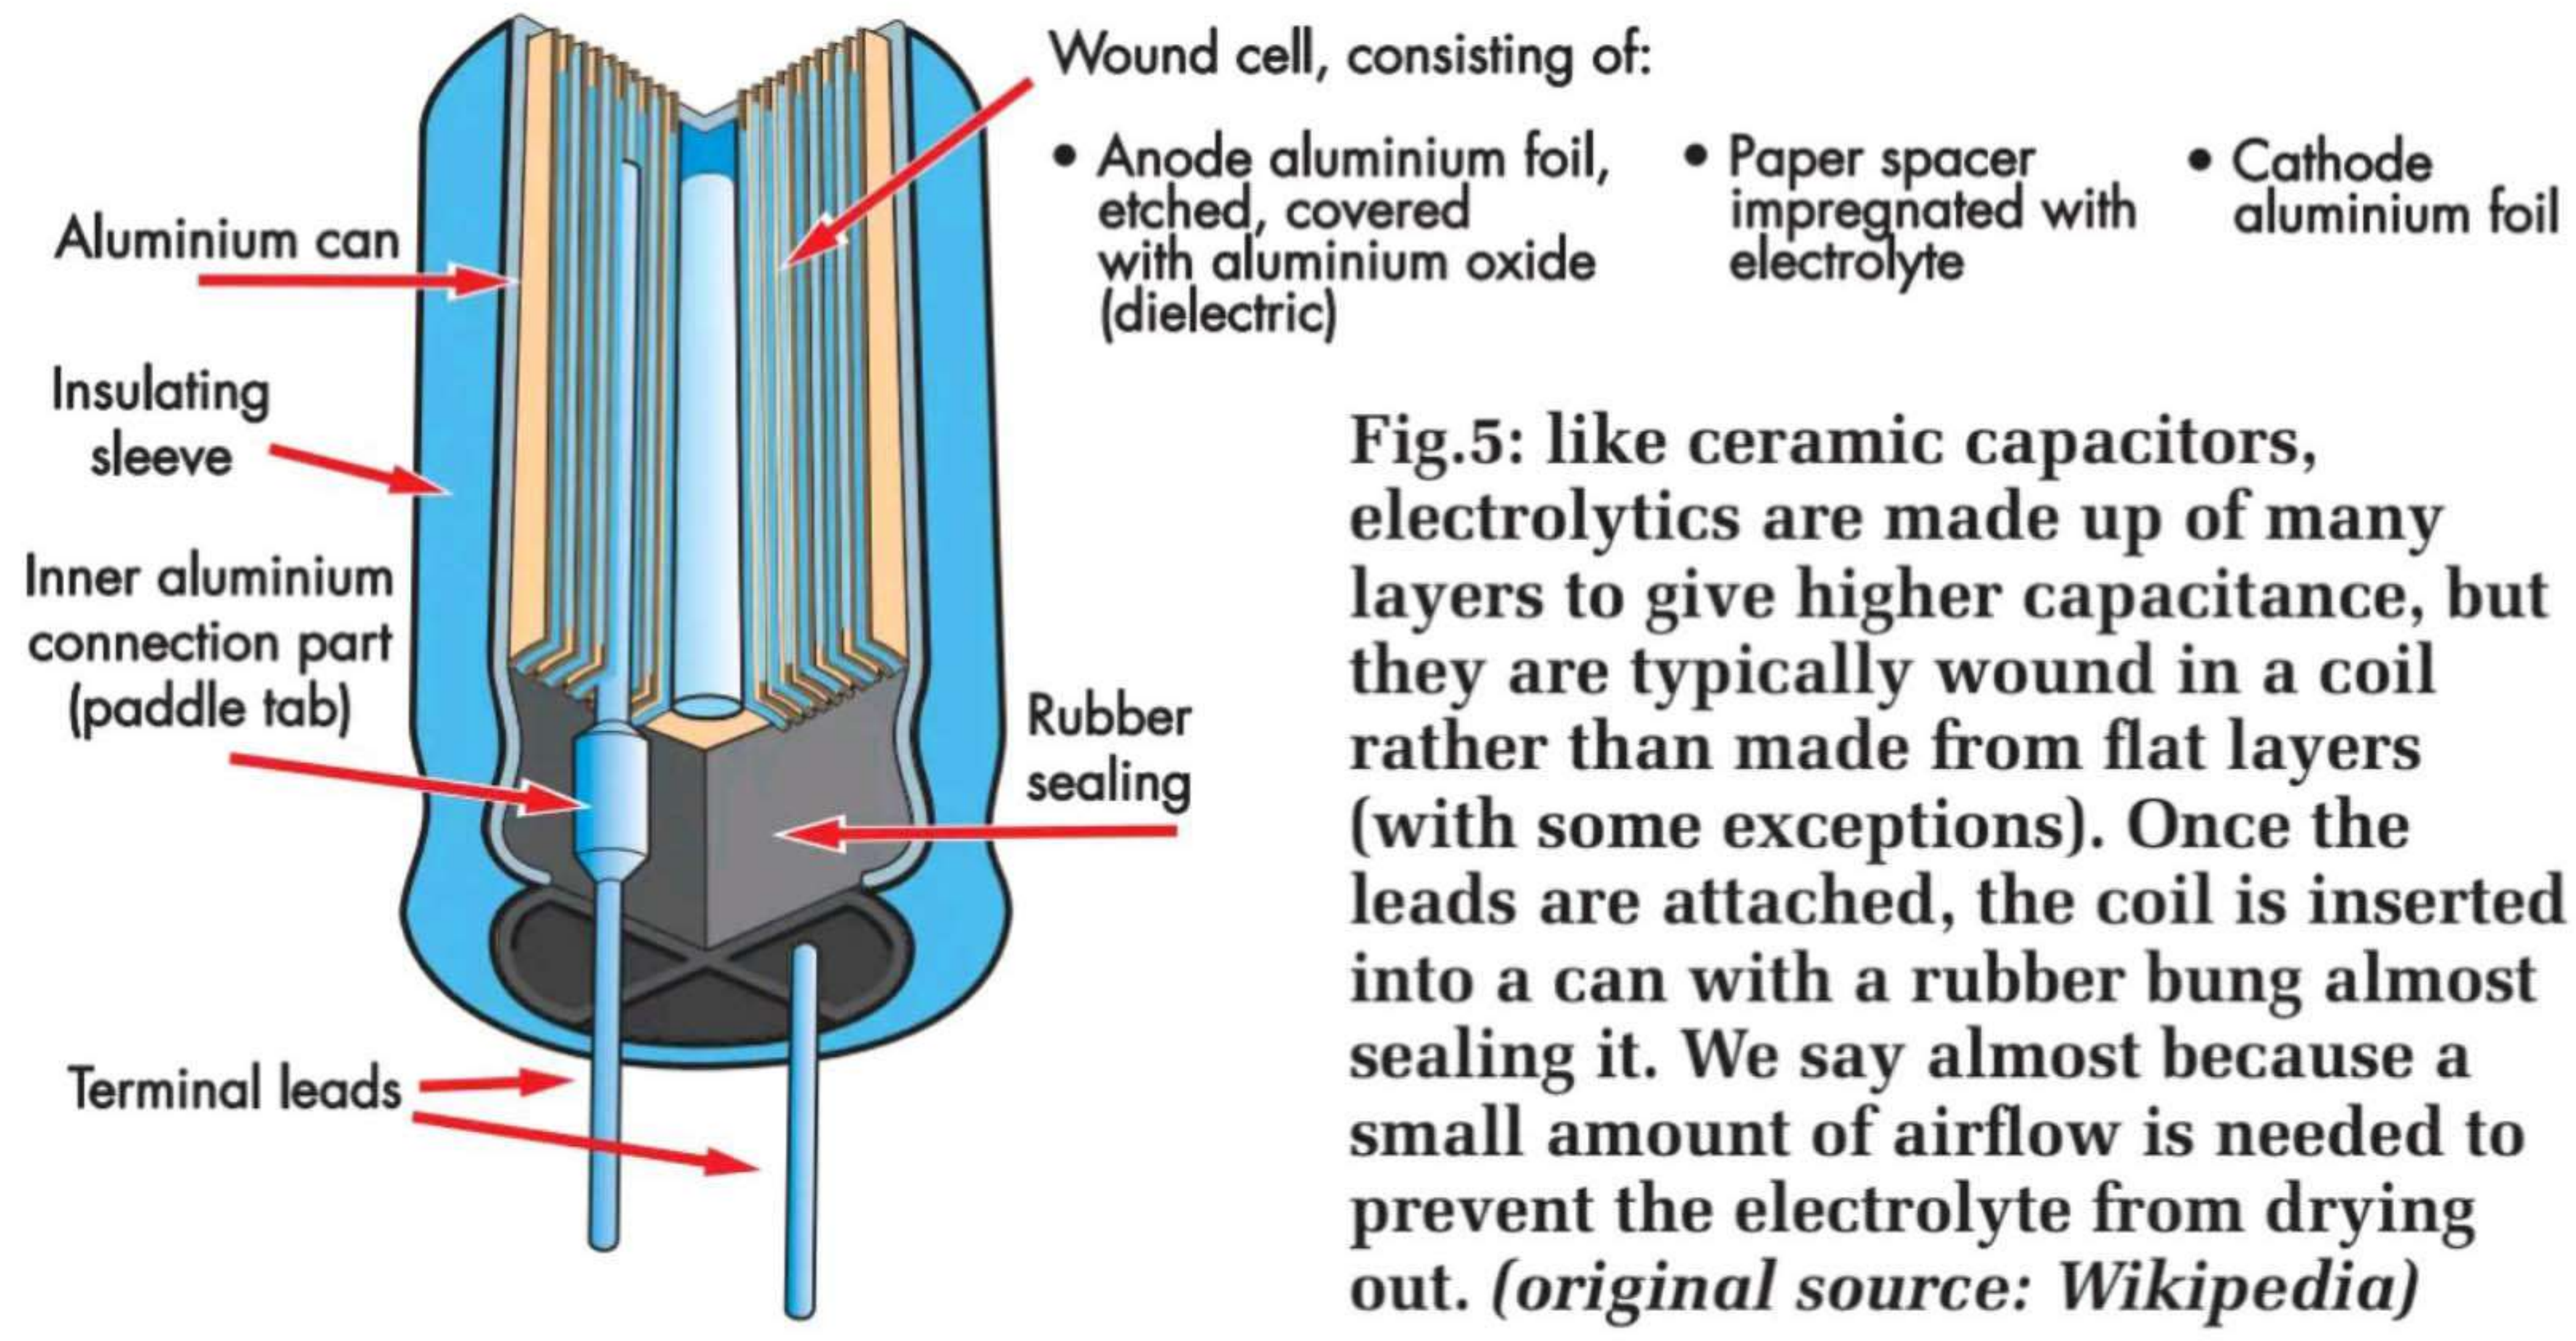
Wound cell, consisting of:
• Anode aluminium foil, etched, covered with aluminium oxide (dielectric)
• Paper spacer impregnated with electrolyte
• Cathode aluminium foil
Aluminium can
Insulating sleeve
Inner aluminium connection part (paddle tab)
Rubber sealing
Terminal leads

Fig.5: like ceramic capacitors, electrolytics are made up of many layers to give higher capacitance, but they are typically wound in a coil rather than made from flat layers (with some exceptions). Once the leads are attached, the coil is inserted into a can with a rubber bung almost sealing it. We say almost because a small amount of airflow is needed to prevent the electrolyte from drying out. *(original source: Wikipedia)*

On the other hand, Y5V ceramics are very compact for a given capacitance value and voltage rating, but they have a very poor initial tolerance, and their capacitance is reduced dramatically at elevated temperatures and higher applied voltages. The benefit is that they are quite cheap to produce.

Dielectrics like X5R and X7R are in between those two; they are larger than Y5V types for a given capacitance and voltage, but smaller than NP0/C0G. Similarly, their tolerances are intermediate, as are their production costs. Therefore, these capacitors are very popular as a reasonable 'middle ground'.

Note that NP0/C0G ceramics are almost ideal capacitors. They have very stable capacitance values over temperature and voltage, low ESL (equivalent series inductance) and ESR (equivalent series resistance), a low dissipation factor and excellent linearity. Their only real disadvantages are a low maximum capacitance value (due to their relatively high volume) and high cost.

We'll describe the meanings of those performance parameters in some detail later in this article.

If the two/three-letter schemes given above look like gobbledygook to you, that might be because there are multiple different naming/categorisation schemes.

The most common schemes are from the EIA, which consist of a letter, a number and a letter. But they don't always mean the same things.

For the most common type of ceramic capacitors (Class 2), the first letter and following number refer to the minimum and maximum temperature range, while the third letter gives the tolerance of the capacitance over this range – see Table 1.

Better-performing capacitors are the ones with a smaller capacitance change over a broader range, like those with an X7R dielectric.

Electrolytic capacitors

In an electrolytic capacitor, one plate is a metal foil while the other 'plate' is a conductive liquid or gel solution, known as the electrolyte.

The metal foil is etched to increase its surface area greatly, and the liquid or gel is in intimate contact with this foil, separated only by a very thin oxide layer. Therefore, electrolytic capacitors have very high capacitance values for their volume (see Figs.4 & 5).

For this reason, "electros" are typically used for 'bulk' bypassing or filtering applications. Their asymmetry, and the fact that the oxide layer is formed by electrons flowing from one 'plate' to the other, means that they are generally polarised. So one of their leads must always be at a higher voltage than the other (although there are ways around this – described later).

Electrolytic capacitors typically range from a bit under 1µF up to 100,000µF or more, with voltage ratings from a few volts to several hundred volts.

Traditionally, the metal foil used was aluminium. However, other metals can be used, giving certain performance advantages. For example, tantalum, while more expensive, generally results in a capacitor which can handle more current and with a lower ESR (see Figs.6 & 7).

Further refinements to the electrolytic capacitor came with the discovery that an organic polymer gel could be used as the electrolyte, giving a roughly ten-fold decrease in overall ESR (Figs.8 & 9).

As such, organic polymer ('solid') aluminium electrolytic capacitors outperform standard tantalum capacitors, and solid tantalum capacitors perform even better again, approaching that of some ceramics but with a better capacitance-to-volume ratio (see Fig.10).

Other, more exotic types of electros are hybrid polymer electros (which combine liquid and polymer electrolytics) and niobium polymer electros; niobium is cheaper than tantalum but performs similarly.

Tantalum and solid electrolytics tend to occupy the space between traditional electros, which are still

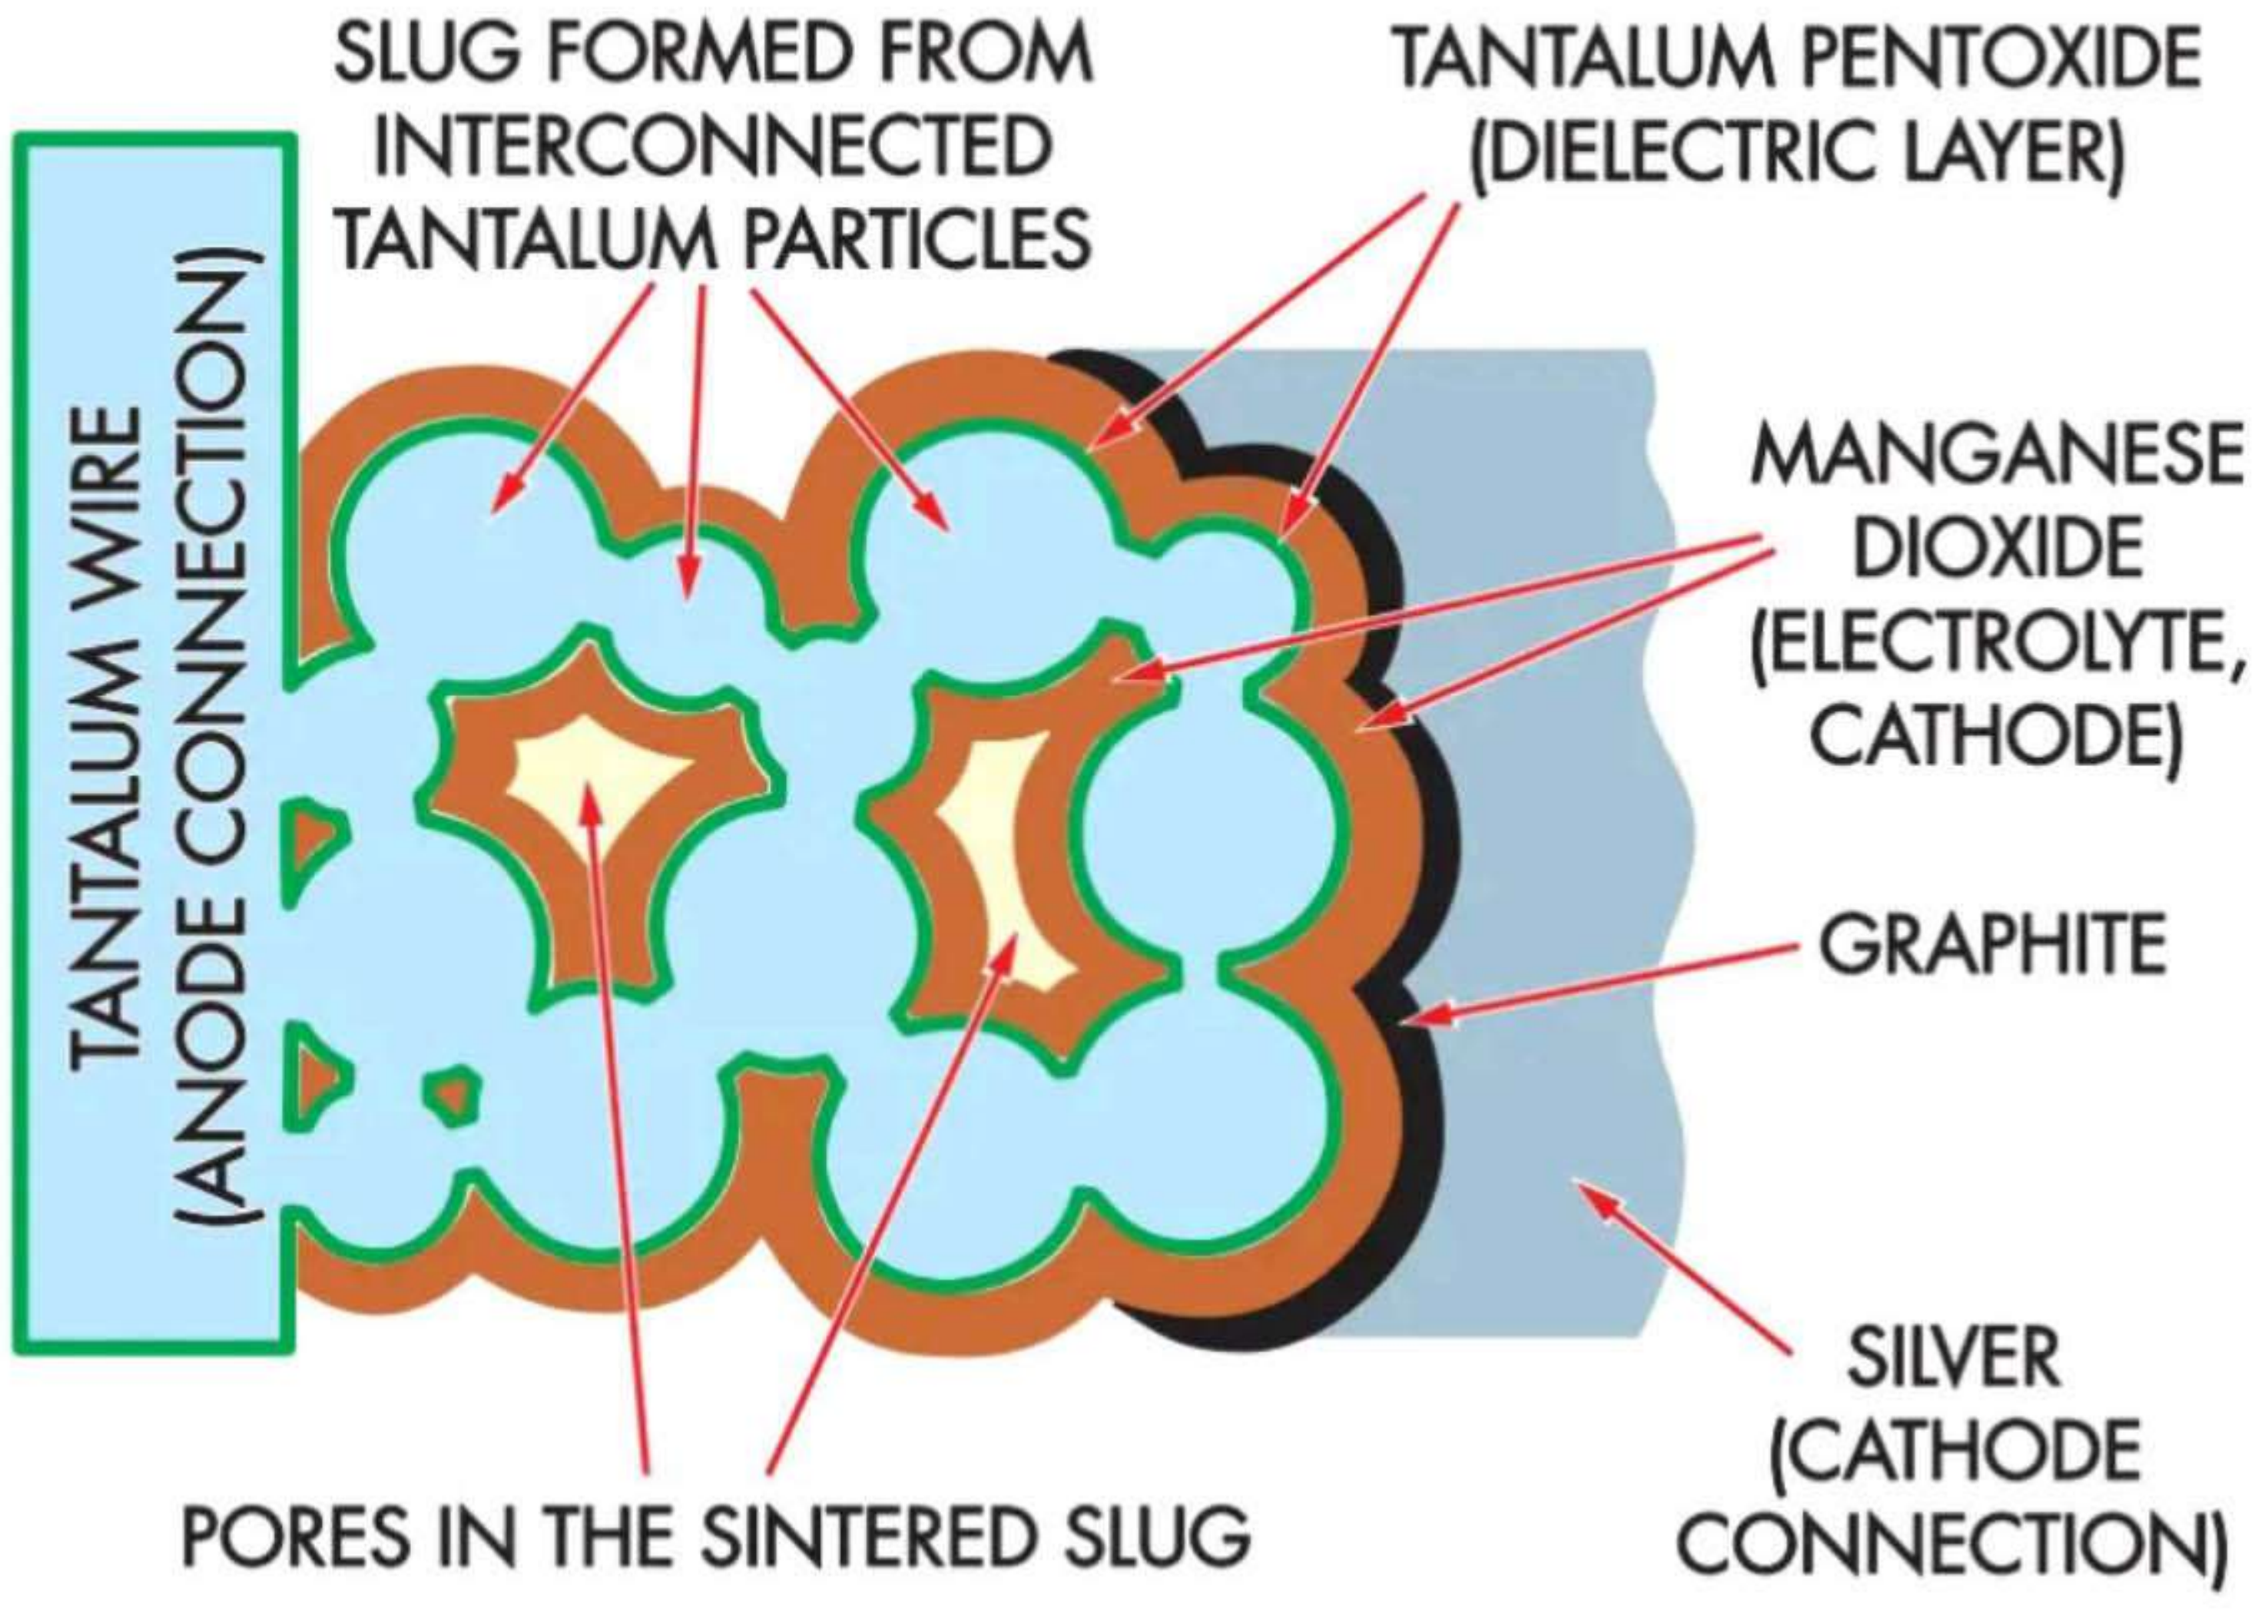
SLUG FORMED FROM INTERCONNECTED TANTALUM PARTICLES
TANTALUM PENTOXIDE (DIELECTRIC LAYER)
TANTALUM WIRE (ANODE CONNECTION)
MANGANESE DIOXIDE (ELECTROLYTE, CATHODE)
GRAPHITE
SILVER (CATHODE CONNECTION)
PORES IN THE SINTERED SLUG

Fig.6: while tantalum electrolytics work on the same principle as aluminium types, the construction method is necessarily quite different due to the properties of tantalum. Tantalum particles are sintered to form a porous slug, which is then immersed in a manganese dioxide electrolyte. Graphite and silver in contact with the electrolyte form the cathode connection, while tantalum wire acts as the anode. ***(original source: Wikipedia)***

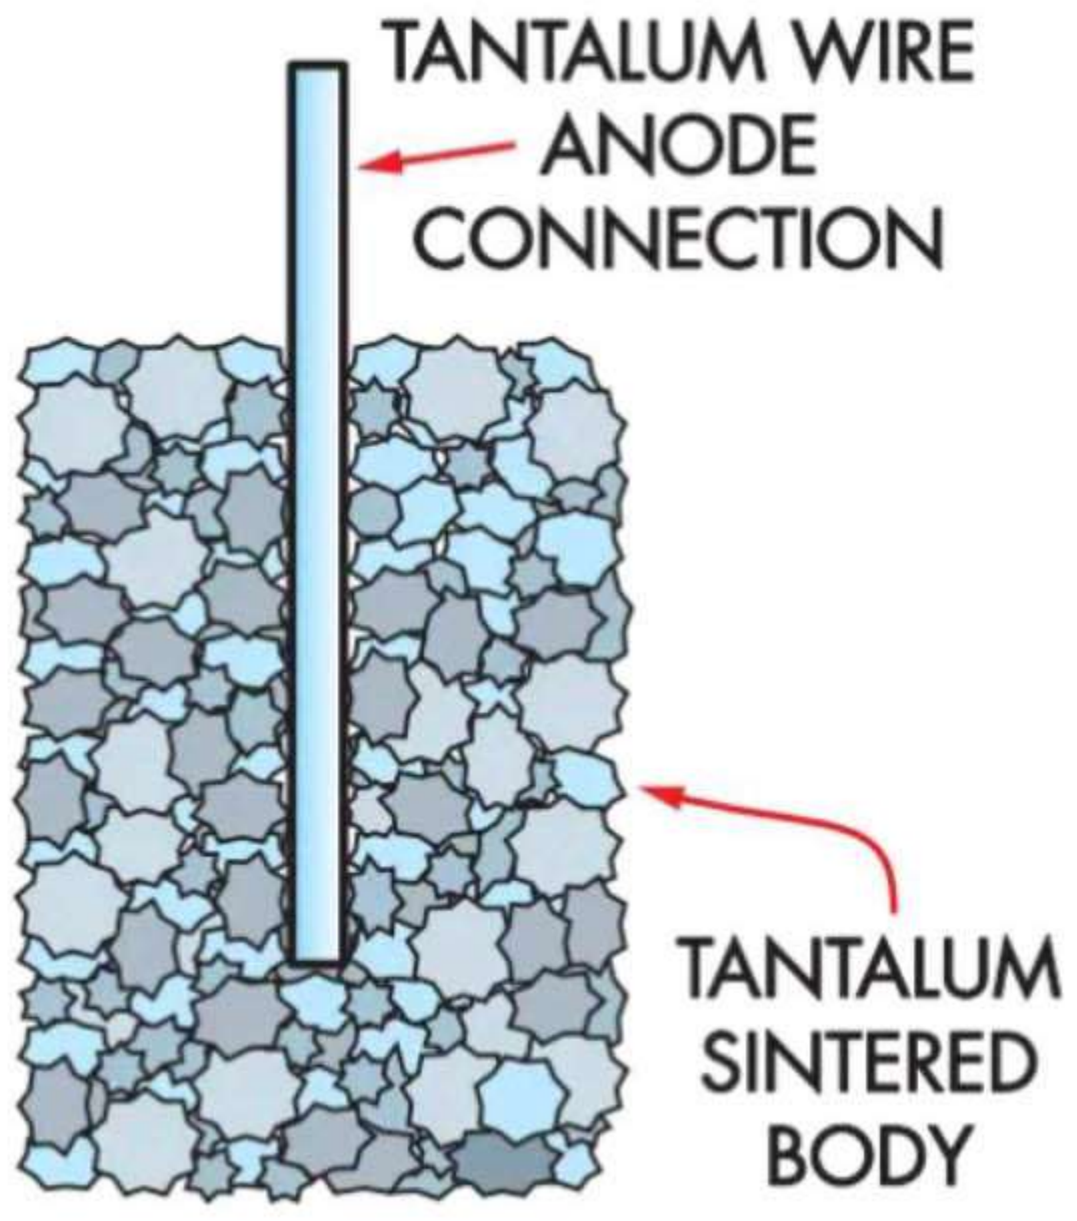
TANTALUM WIRE ANODE CONNECTION
TANTALUM SINTERED BODY

Fig.7: this gives you an idea of how the tantalum particles are sintered and attached to a tantalum wire lead to form the body of the capacitor. ***(original source: Wikipedia)***

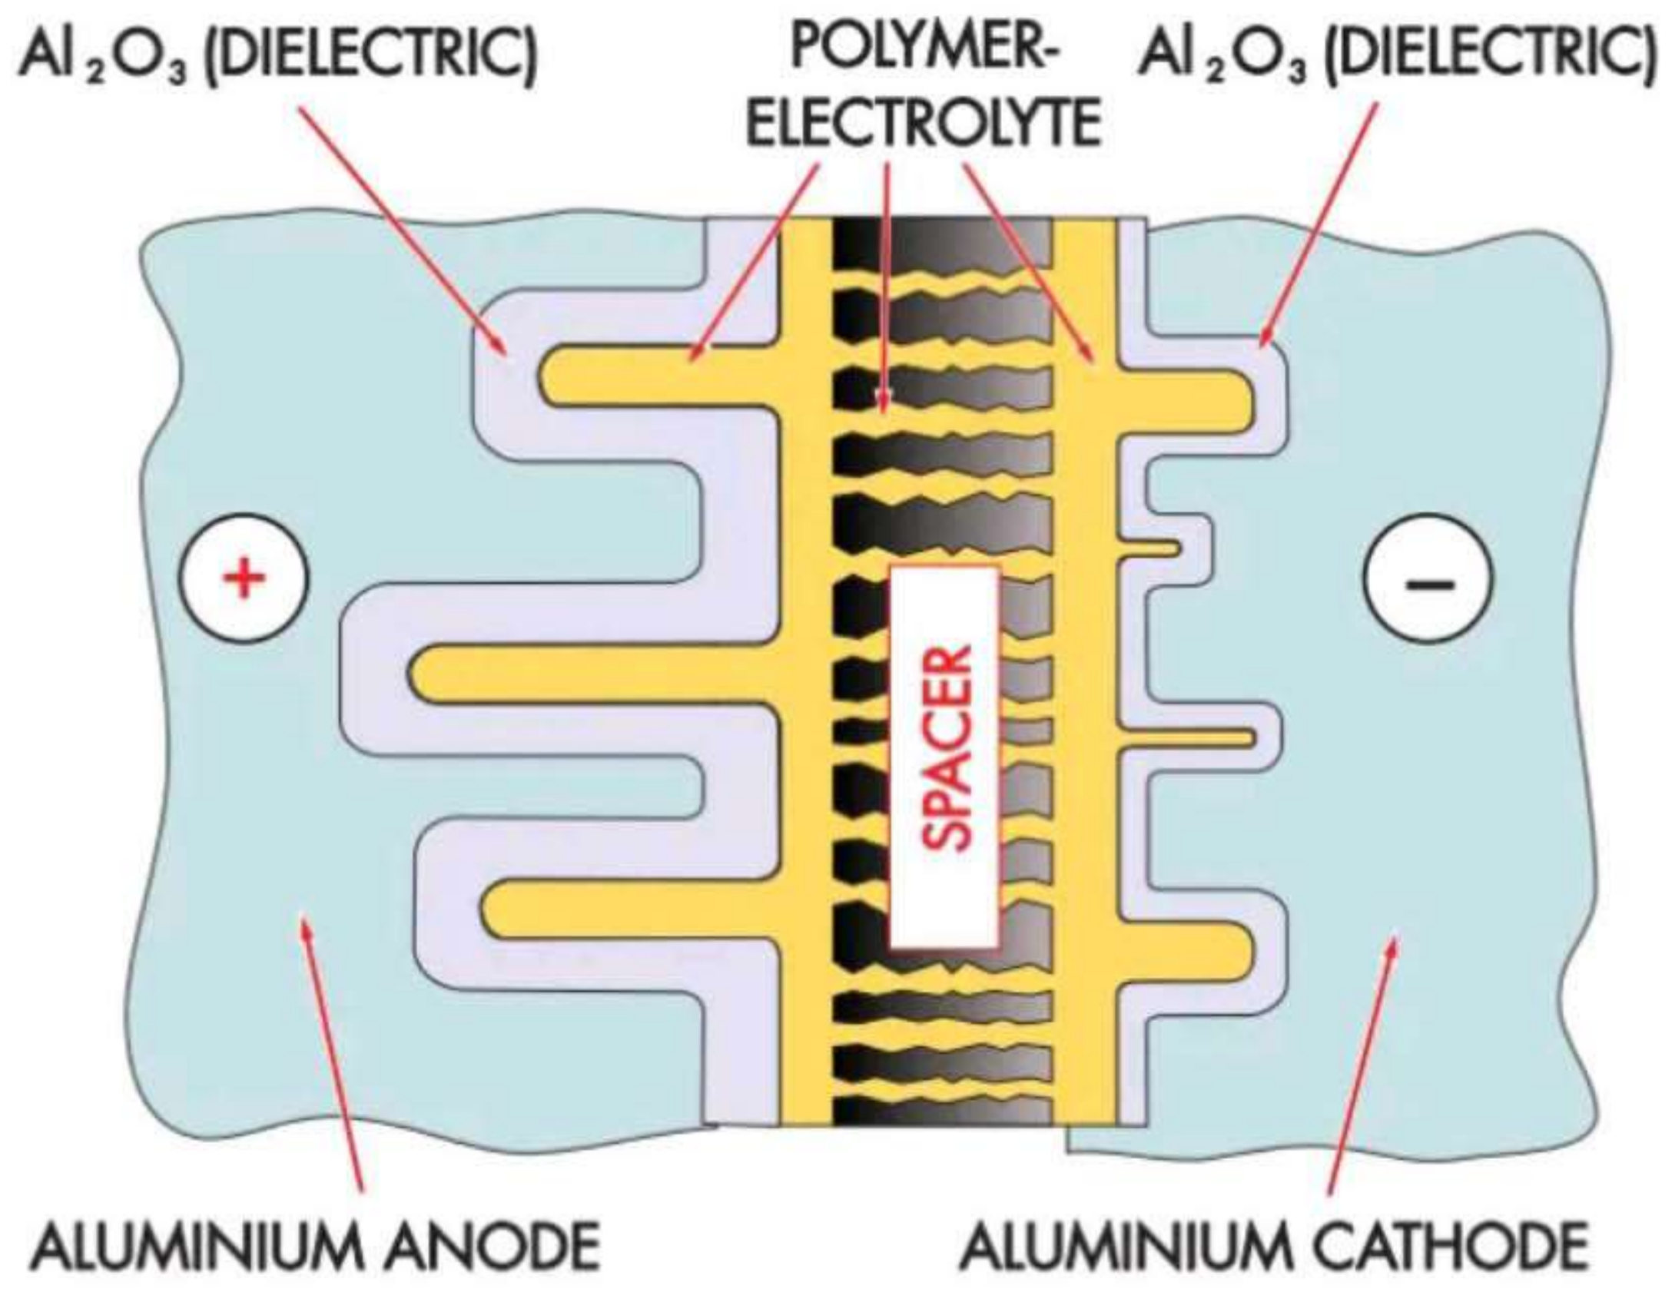
Al₂O₃ (DIELECTRIC)
POLYMER-ELECTROLYTE
Al₂O₃ (DIELECTRIC)
+
SPACER
–
ALUMINIUM ANODE
ALUMINIUM CATHODE

Fig.8: polymer or 'solid' aluminium electrolytic capacitors use an organic polymer (plastic-like) electrolyte which has roughly ten times the conductivity of a wet electrolyte. This allows for more compact electrolytic capacitors with much higher ripple current ratings and lower ESR values. Other techniques like comprehensive lead stitching are often combined with the polymer electrolyte to maximise performance. ***(original source: Wikipedia)***

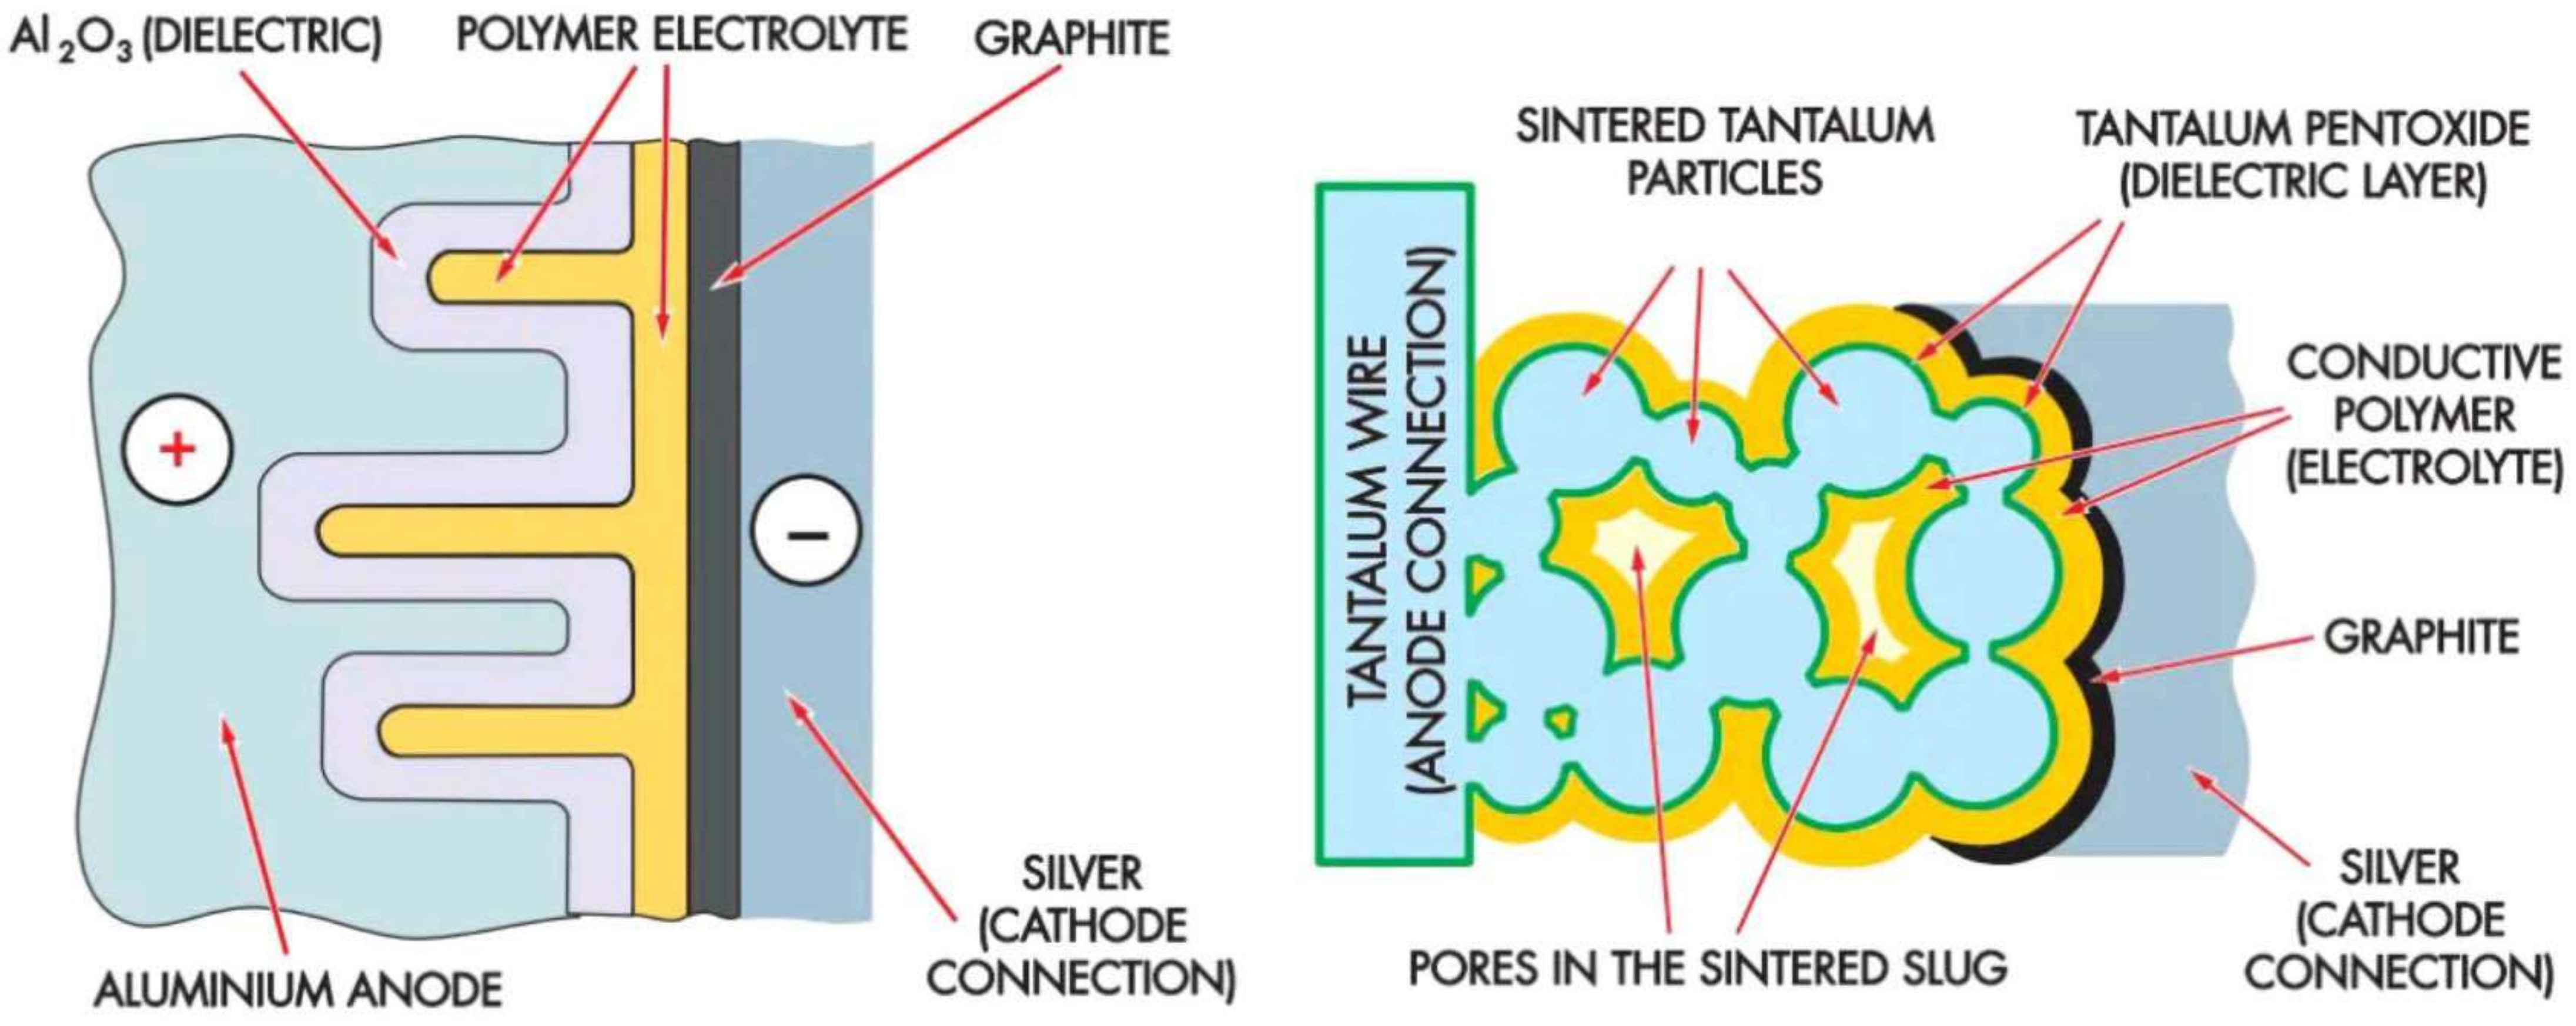
Al₂O₃ (DIELECTRIC)
POLYMER ELECTROLYTE
GRAPHITE
+
–
ALUMINIUM ANODE
SILVER (CATHODE CONNECTION)
SINTERED TANTALUM PARTICLES
TANTALUM PENTOXIDE (DIELECTRIC LAYER)
TANTALUM WIRE (ANODE CONNECTION)
CONDUCTIVE POLYMER (ELECTROLYTE)
GRAPHITE
SILVER (CATHODE CONNECTION)
PORES IN THE SINTERED SLUG

Figs.9(a) & (b): an alternative construction for polymer electrolytic capacitors which uses the same cathode construction as a sintered tantalum capacitor. This halves the number of dielectric layers, significantly increasing capacitance at the cost of a more complicated manufacturing process and more expensive inputs. Polymer tantalum capacitors are made much the same as regular tantalums, just with a different type of electrolyte. ***(original source: Wikipedia)***

A set of electrolytic capacitors ranging from 10µF to 68mF (68,000µF). Note how the can-type electros have a stripe to indicate the negative lead, while the rectangular SMD types have a stripe showing the positive lead.

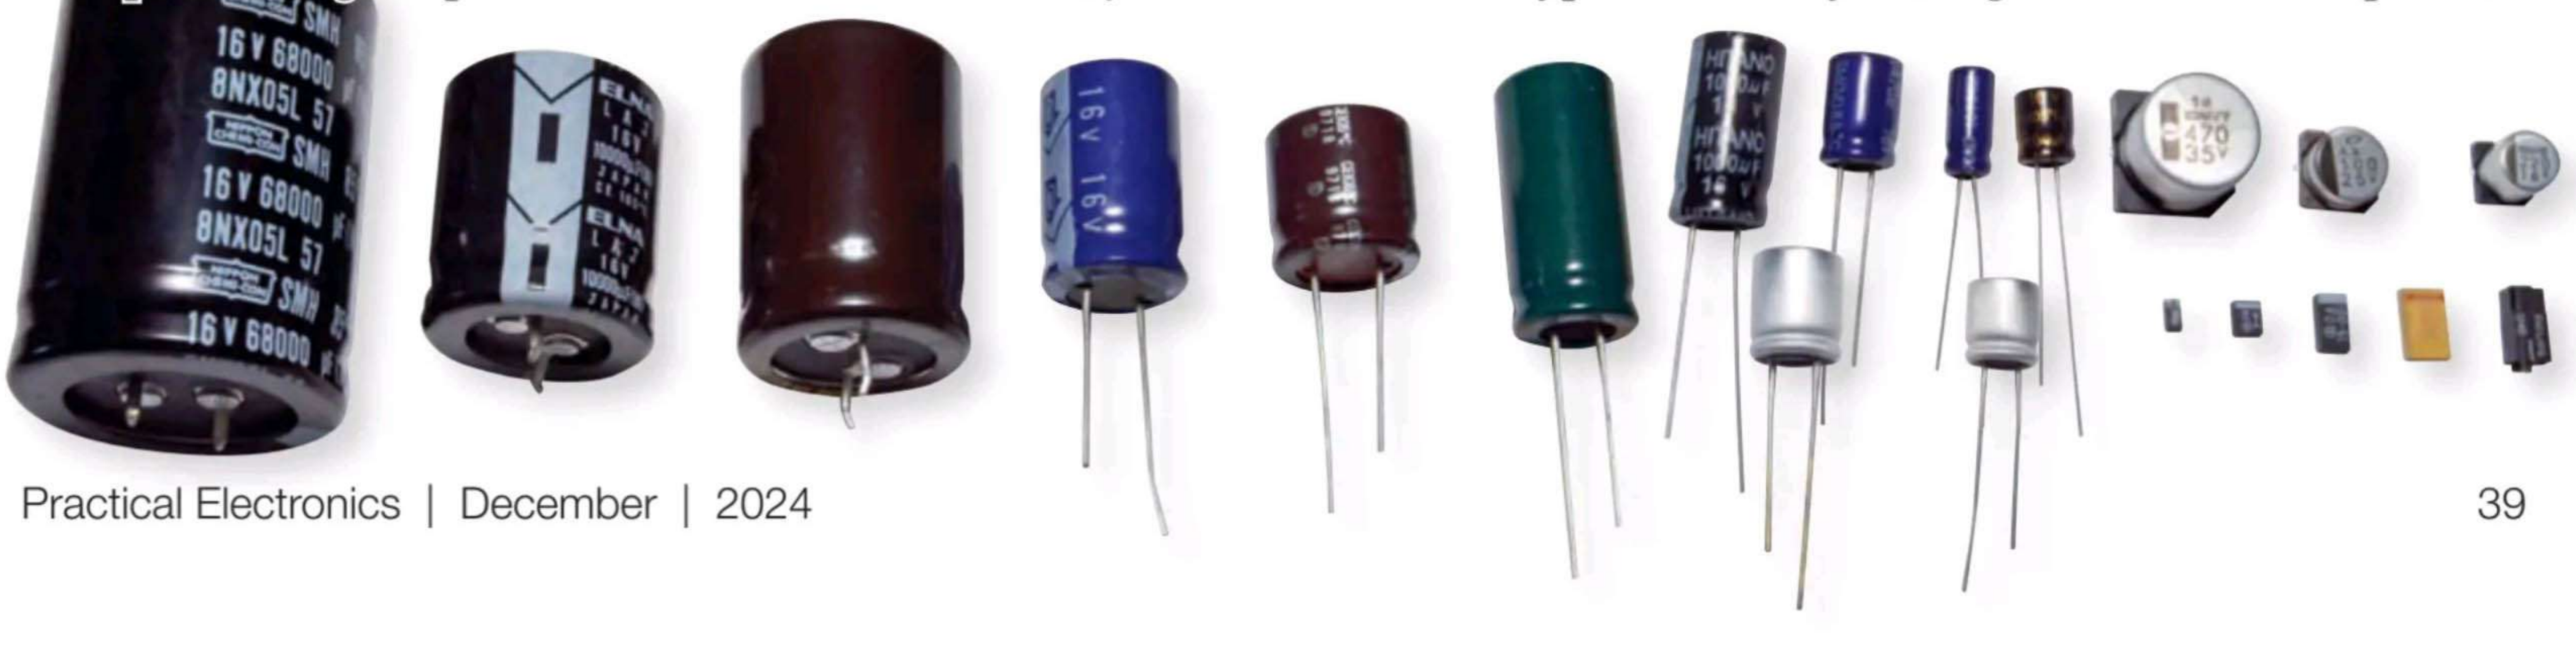

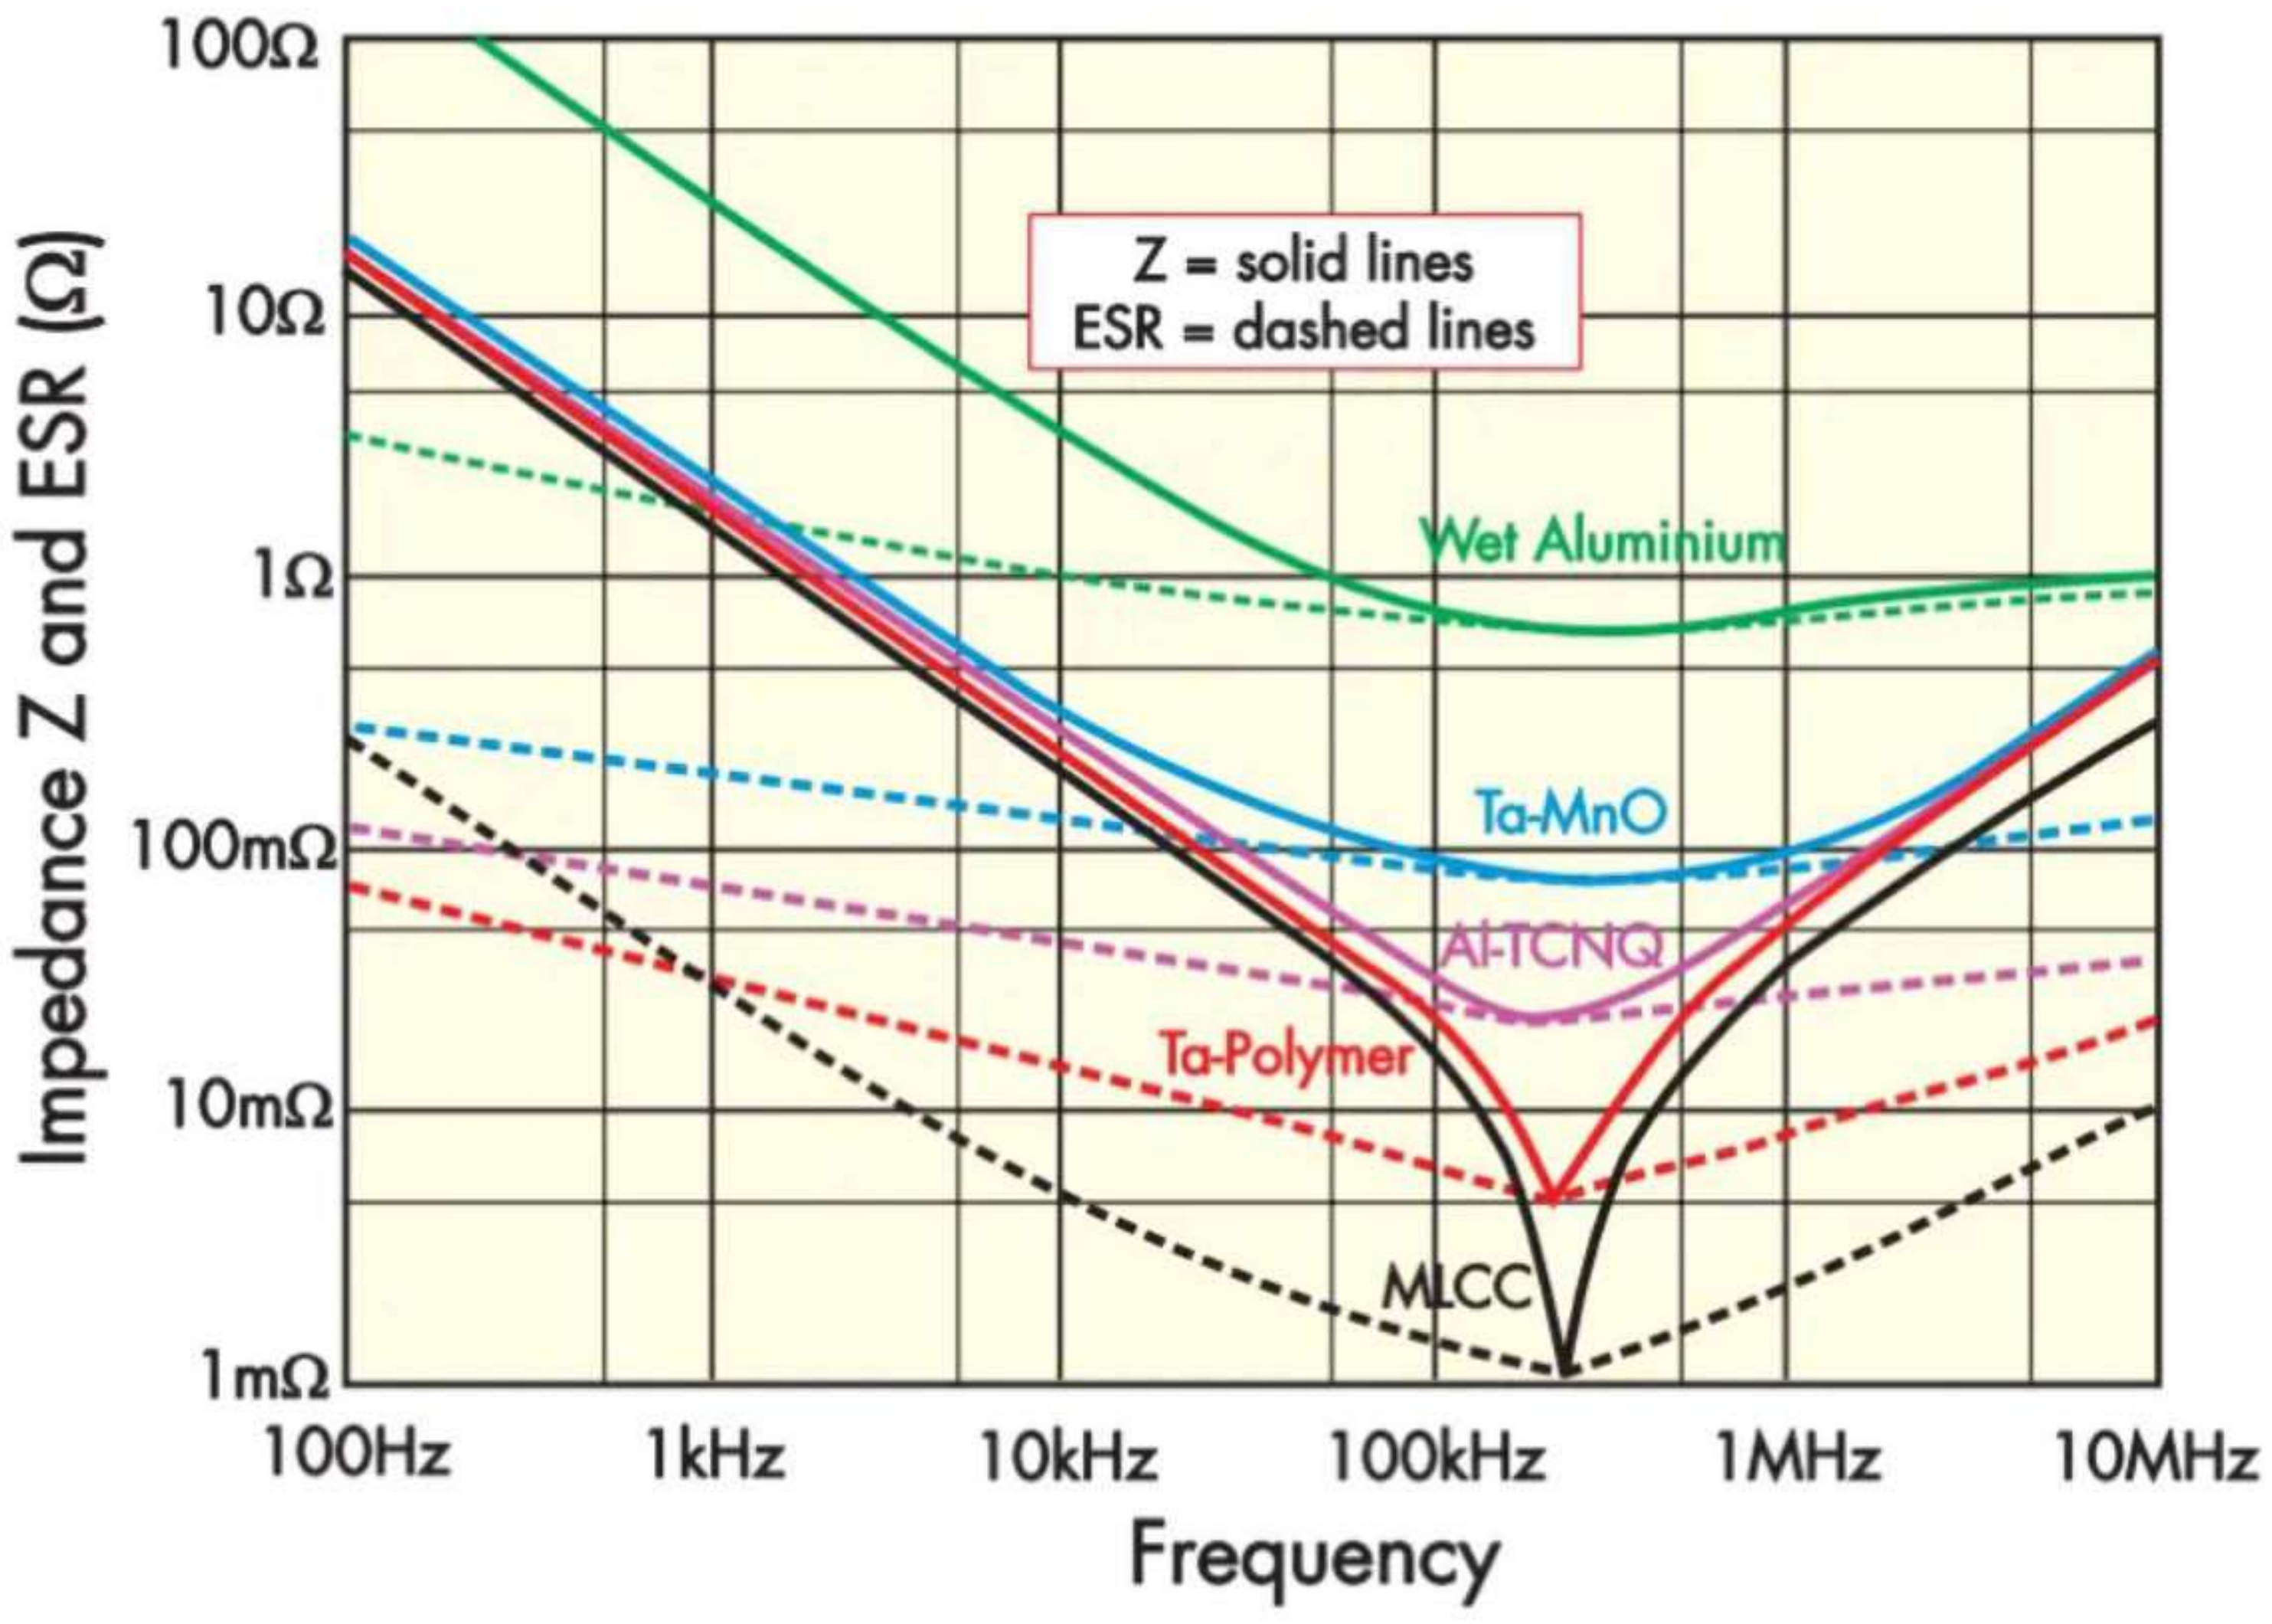

Fig.10: an impedance vs frequency graph comparing four different types of electrolytic capacitor and a multi-layer ceramic capacitor (MLCC) with the same self-resonant frequency. The tantalum-polymer capacitor comes closest to the MLCC in terms of performance at the resonant frequency, while being superior at lower frequencies, likely due to a higher capacitance value. ***(original source: Wikipedia)***

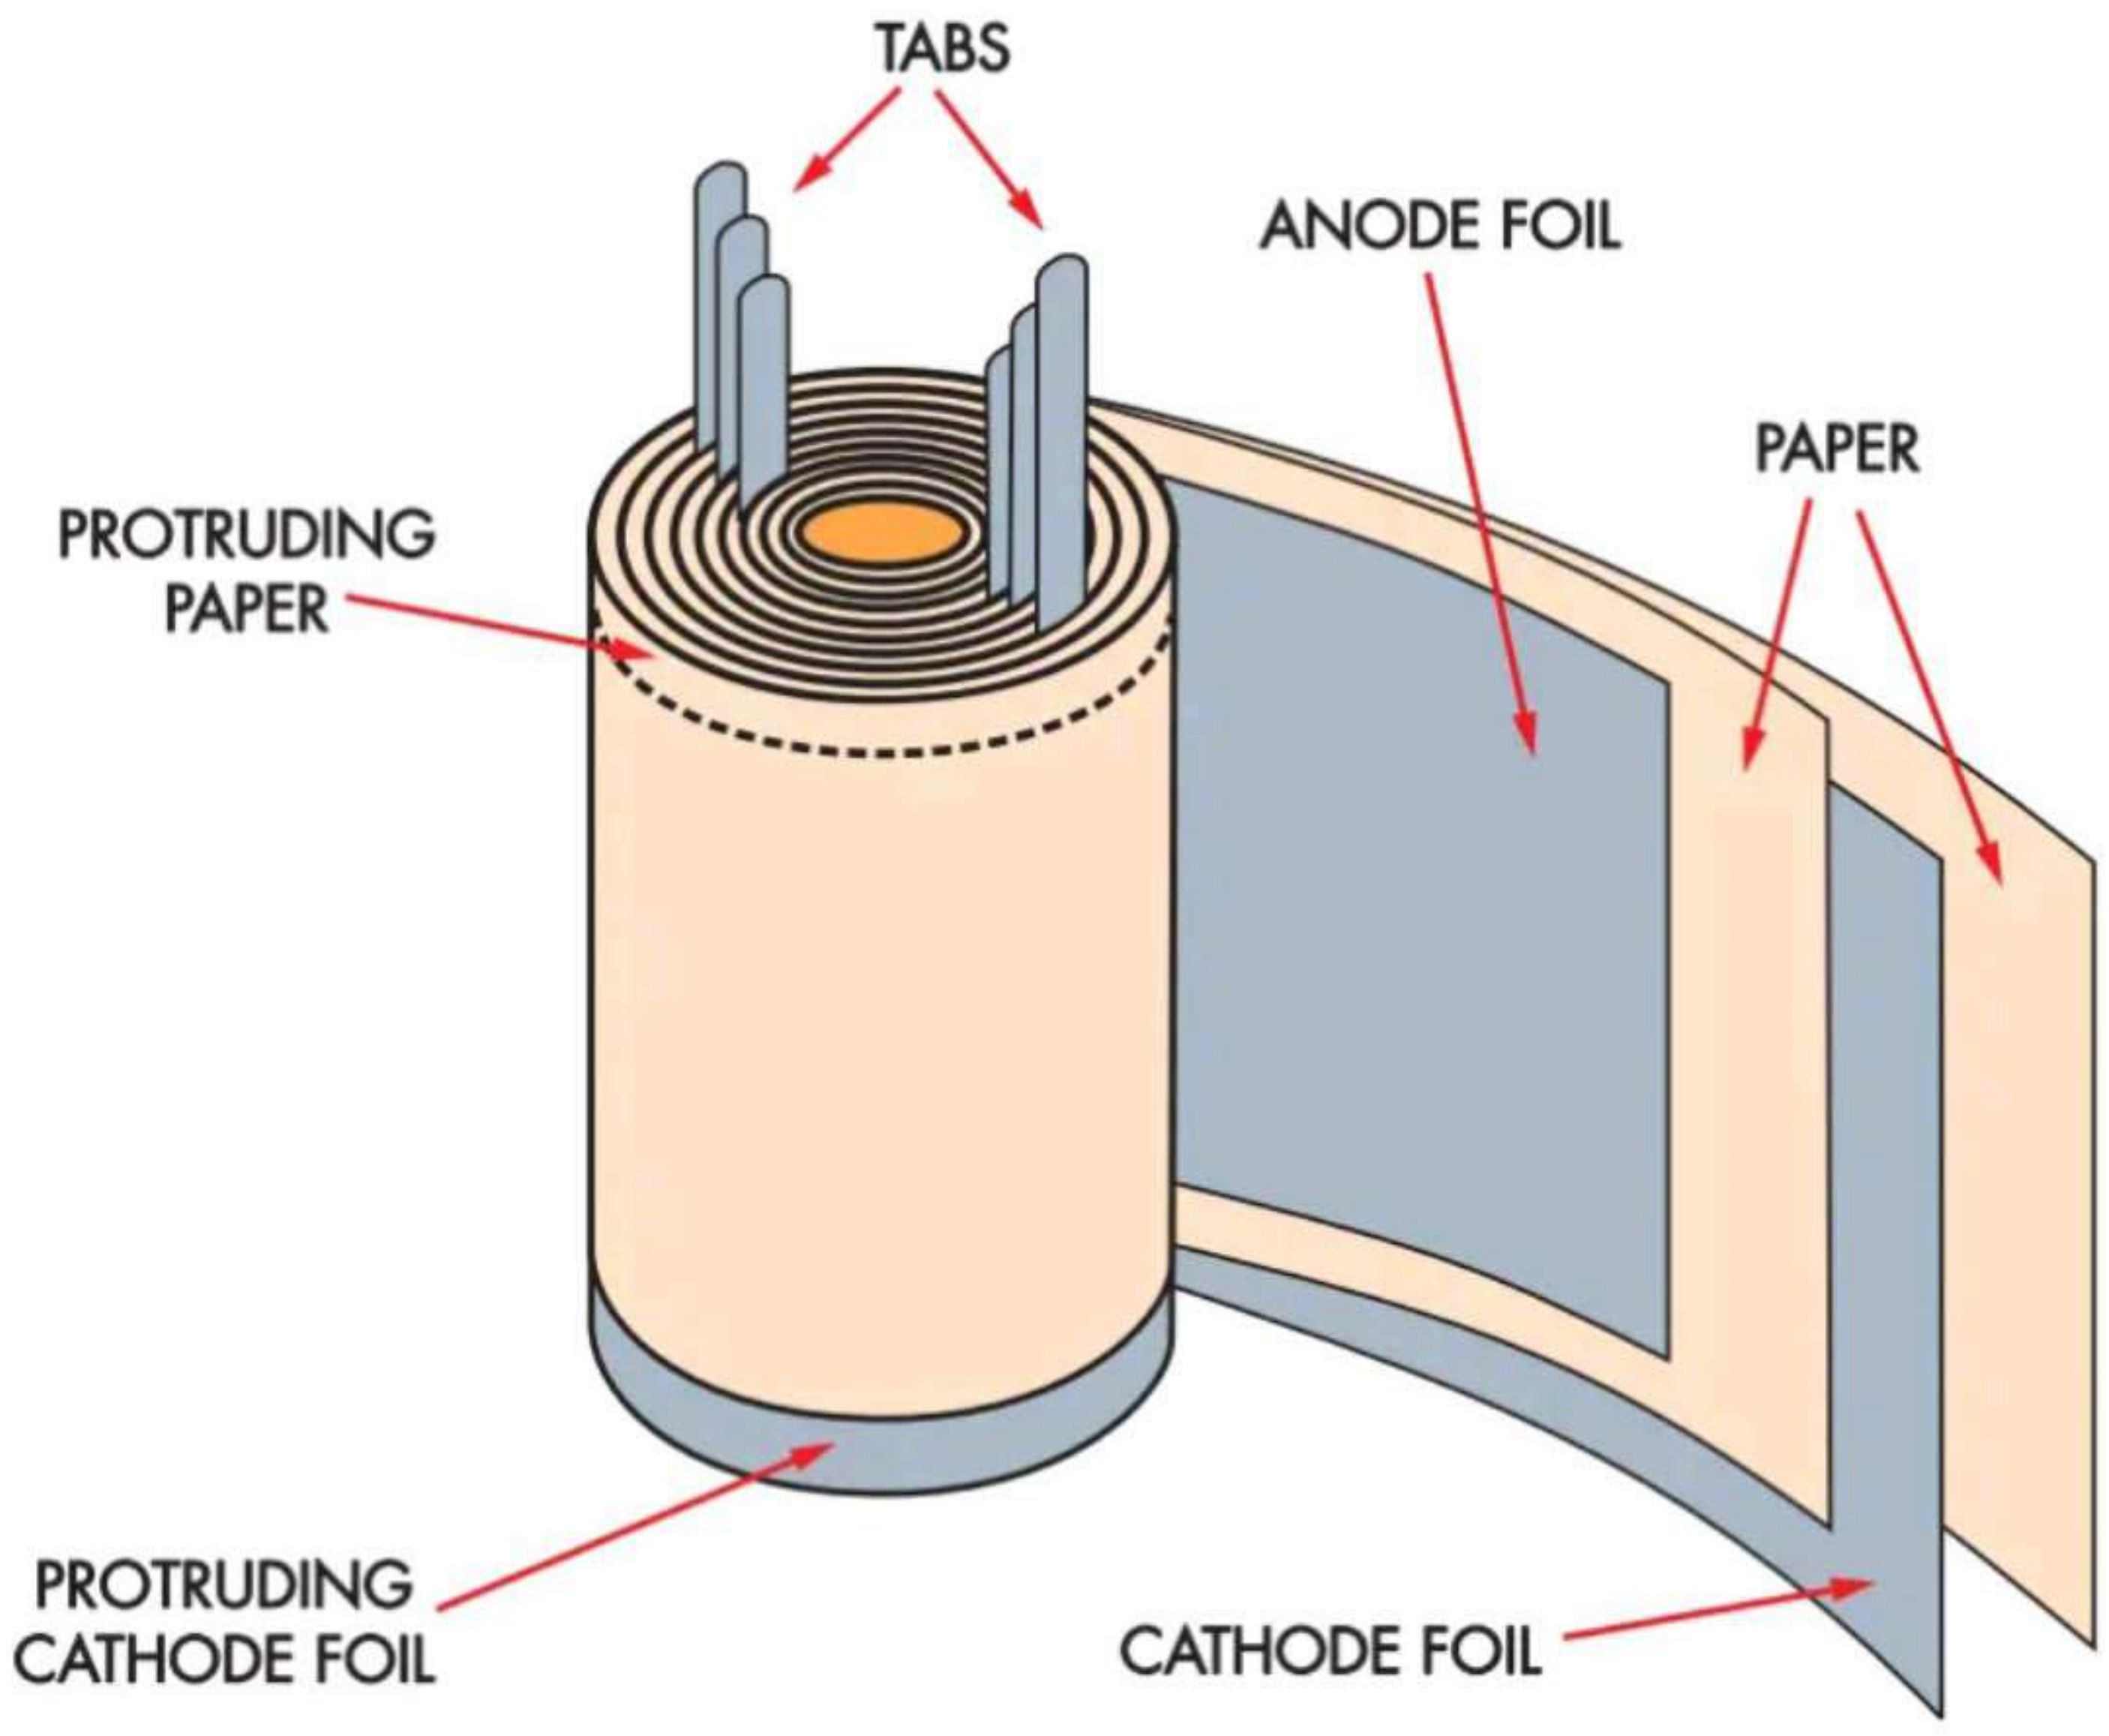

Fig.11: traditional electrolytic capacitors are wound with four layers: two metal foils and two paper separators which are soaked in electrolyte. Note the multiple tabs which ensure that current doesn't have to flow too far to reach any point on the foils. Anode and cathode tabs are lined up together so they can be welded to the appropriate leads. ***(original source: TDK)***

Fig.12: for high-performance (eg, low-ESR) capacitors, the lead tabs are stitched into the aluminium foils, with the metal of the lead tab and foil being joined at multiple points throughout the foil to provide a low-resistance, low-inductance path for current to flow.

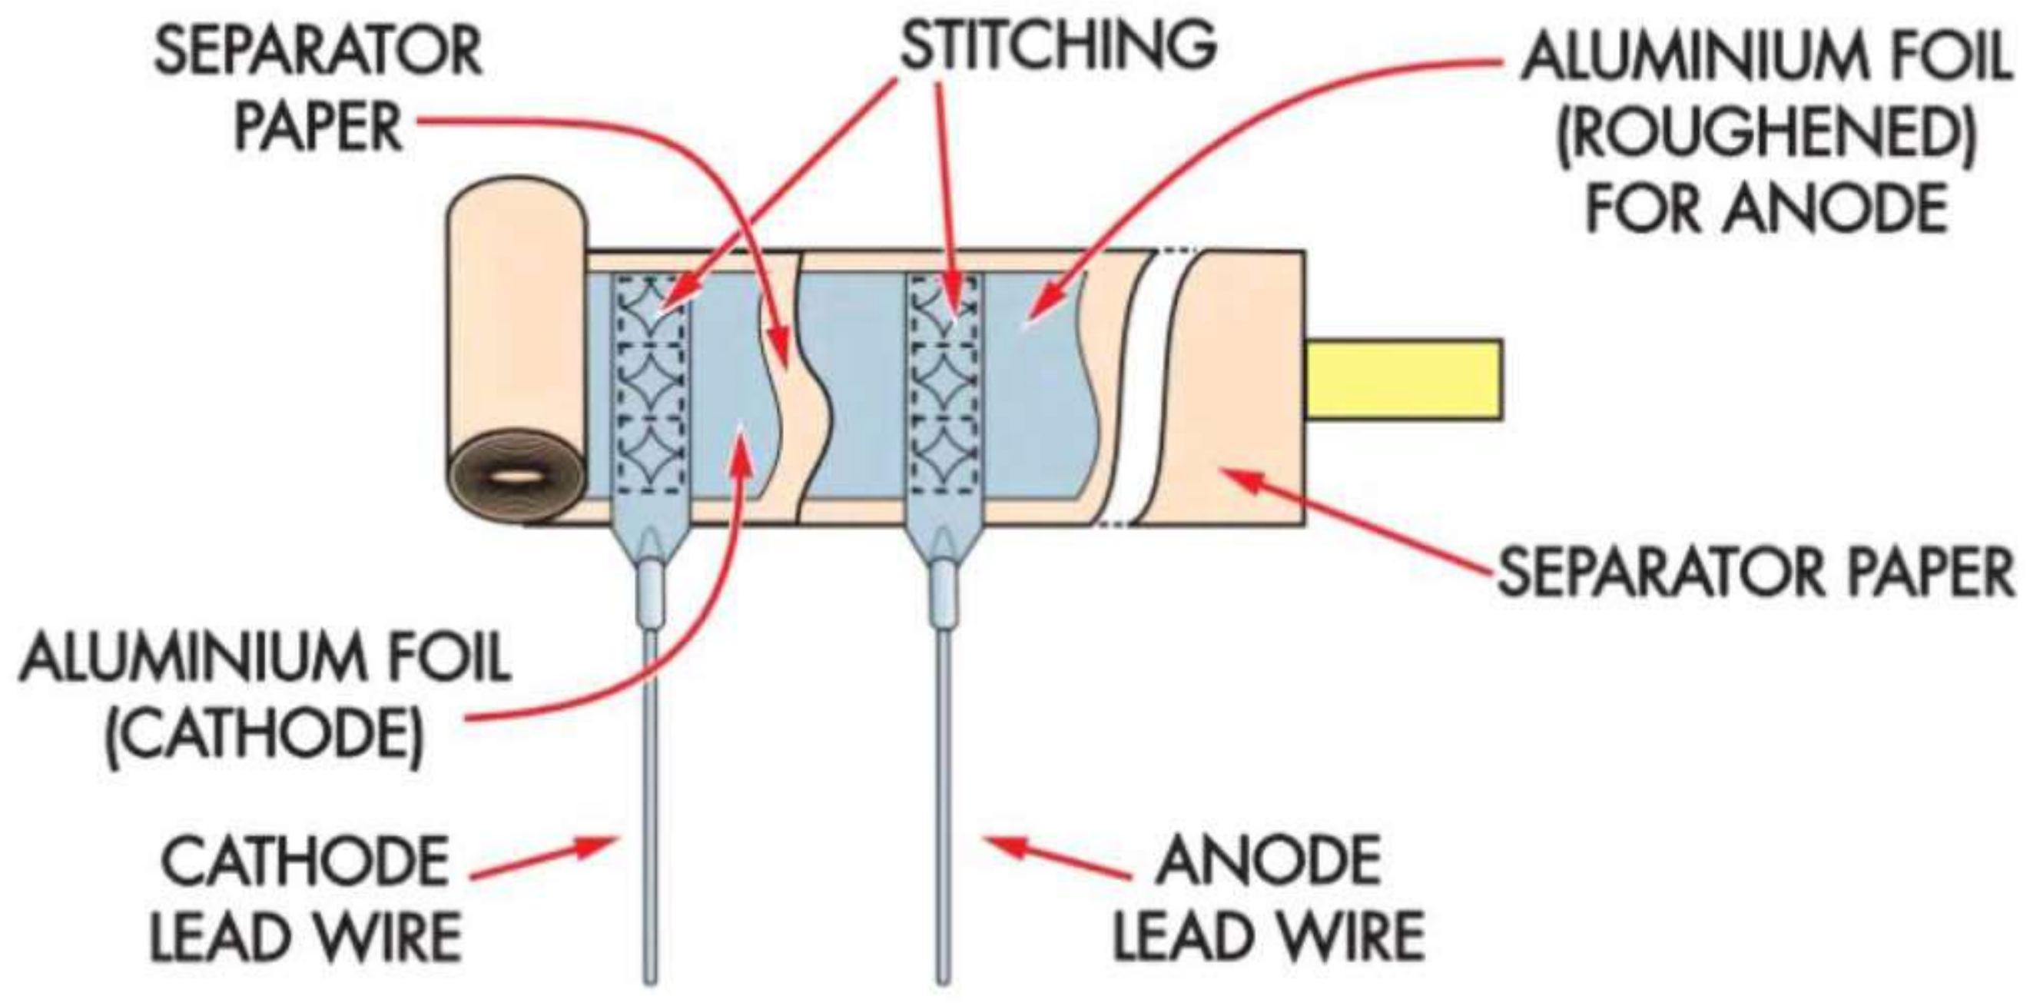

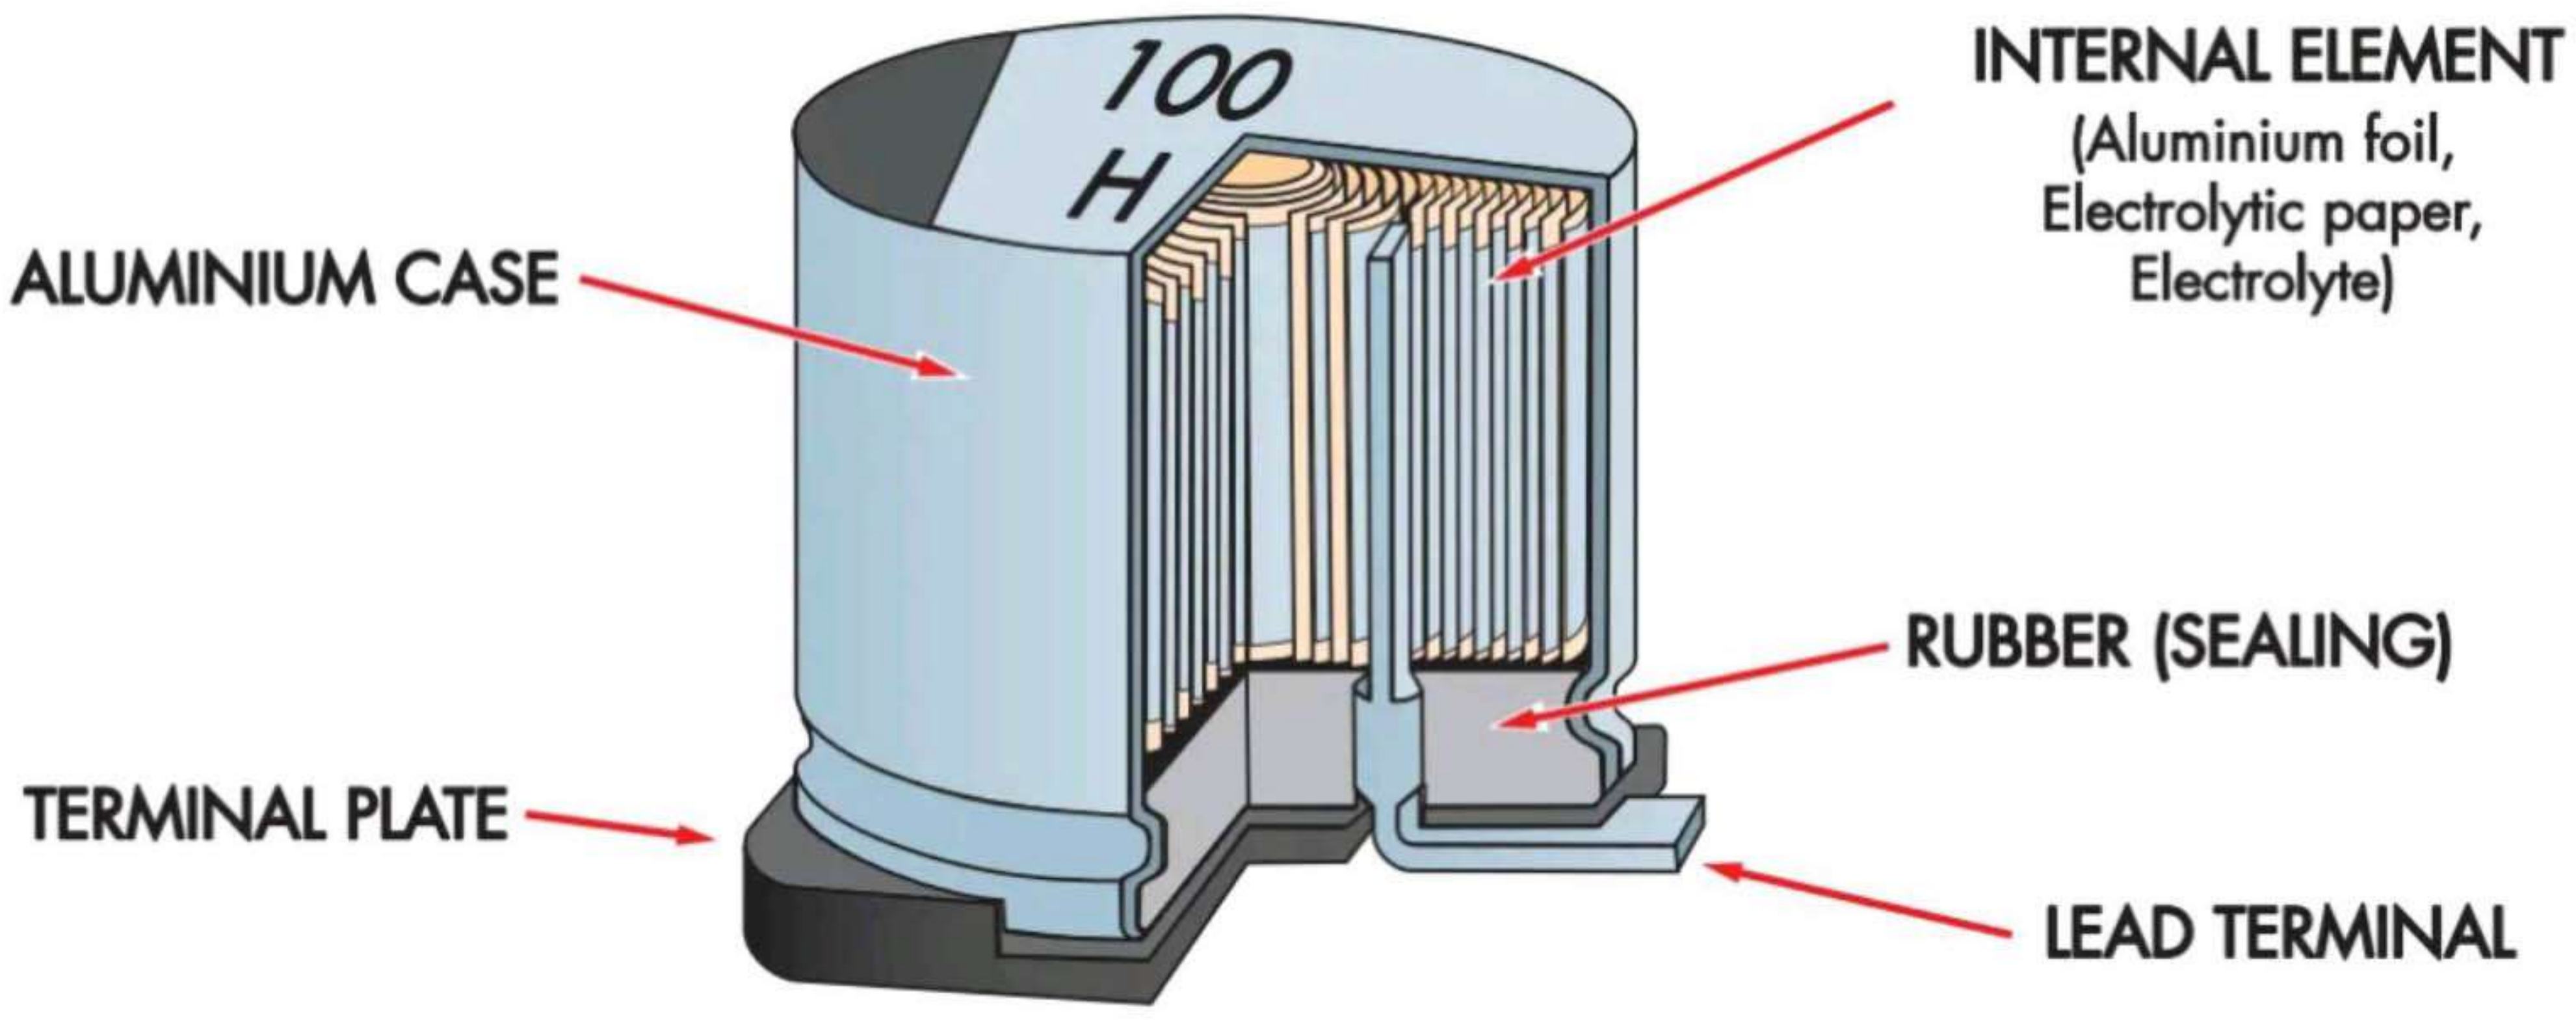

Fig.13: SMD electrolytic capacitors come in different forms, but the can style uses very similar construction to a through-hole radial capacitor. The main differences are that the can sits on a plastic platform with the leads bent horizontally under it, so that the capacitor sits on the PCB and the leads rest on their respective pads. ***(original source: Panasonic)***

widely used for bulk filtering, and ceramics, which are used more for local or high-performance supply bypassing.

For example, tantalum or polymer electros might be used in switch-mode power supply circuitry, where very high pulse current handling and good filtering (low ESR) is critical.

Non-polarised (NP) or bipolar (BP) electrolytic capacitors are effectively two polarised electrolytic capacitors connected back-to-back. You can create an NP electro from two polarised electros by joining either the negative or positive leads together, although the internal construction of an NP/BP may be somewhat different in practice. But the result is the same.

This works because when one of the two capacitors is reverse-biased, it acts a bit like a diode, and the other capacitor does all the work. When the voltage reverses, the capacitors swap roles.

Strangely, you can often get better performance by connecting two polarised electros back-to-back, at a lower cost than a dedicated NP capacitor! This is probably due to economies of scale; polarised electros are made by the squillions while NP capacitors are used in fairly specialised roles.

Another thing to note about electros regarding polarity is that it is safe to apply a reverse polarity voltage long-term as long as it is low (ie, below the threshold where it starts to conduct). This means that polarised electros are quite suitable for use as AC-coupling capacitors even if the polarity of the voltage across them is not known, as long as that voltage never exceeds about ±1.5V DC.

Electrolytic construction

Traditionally, electrolytic capacitors are 'wound'. Two long strips of aluminium foil are chemically etched to increase their surface area, then a strip of paper (or some other porous insulator) soaked in electrolyte is sandwiched between them.

Each conductive strip has one or more tags, for ultimately attaching leads, projecting from one side (see Figs.11-13).

This sandwich is then wound into a roll, with a second paper layer to separate the anode and cathode. With the leads in place, the roll assembly is inserted into a can. More electrolyte is added, and a rubber bung to seal it.

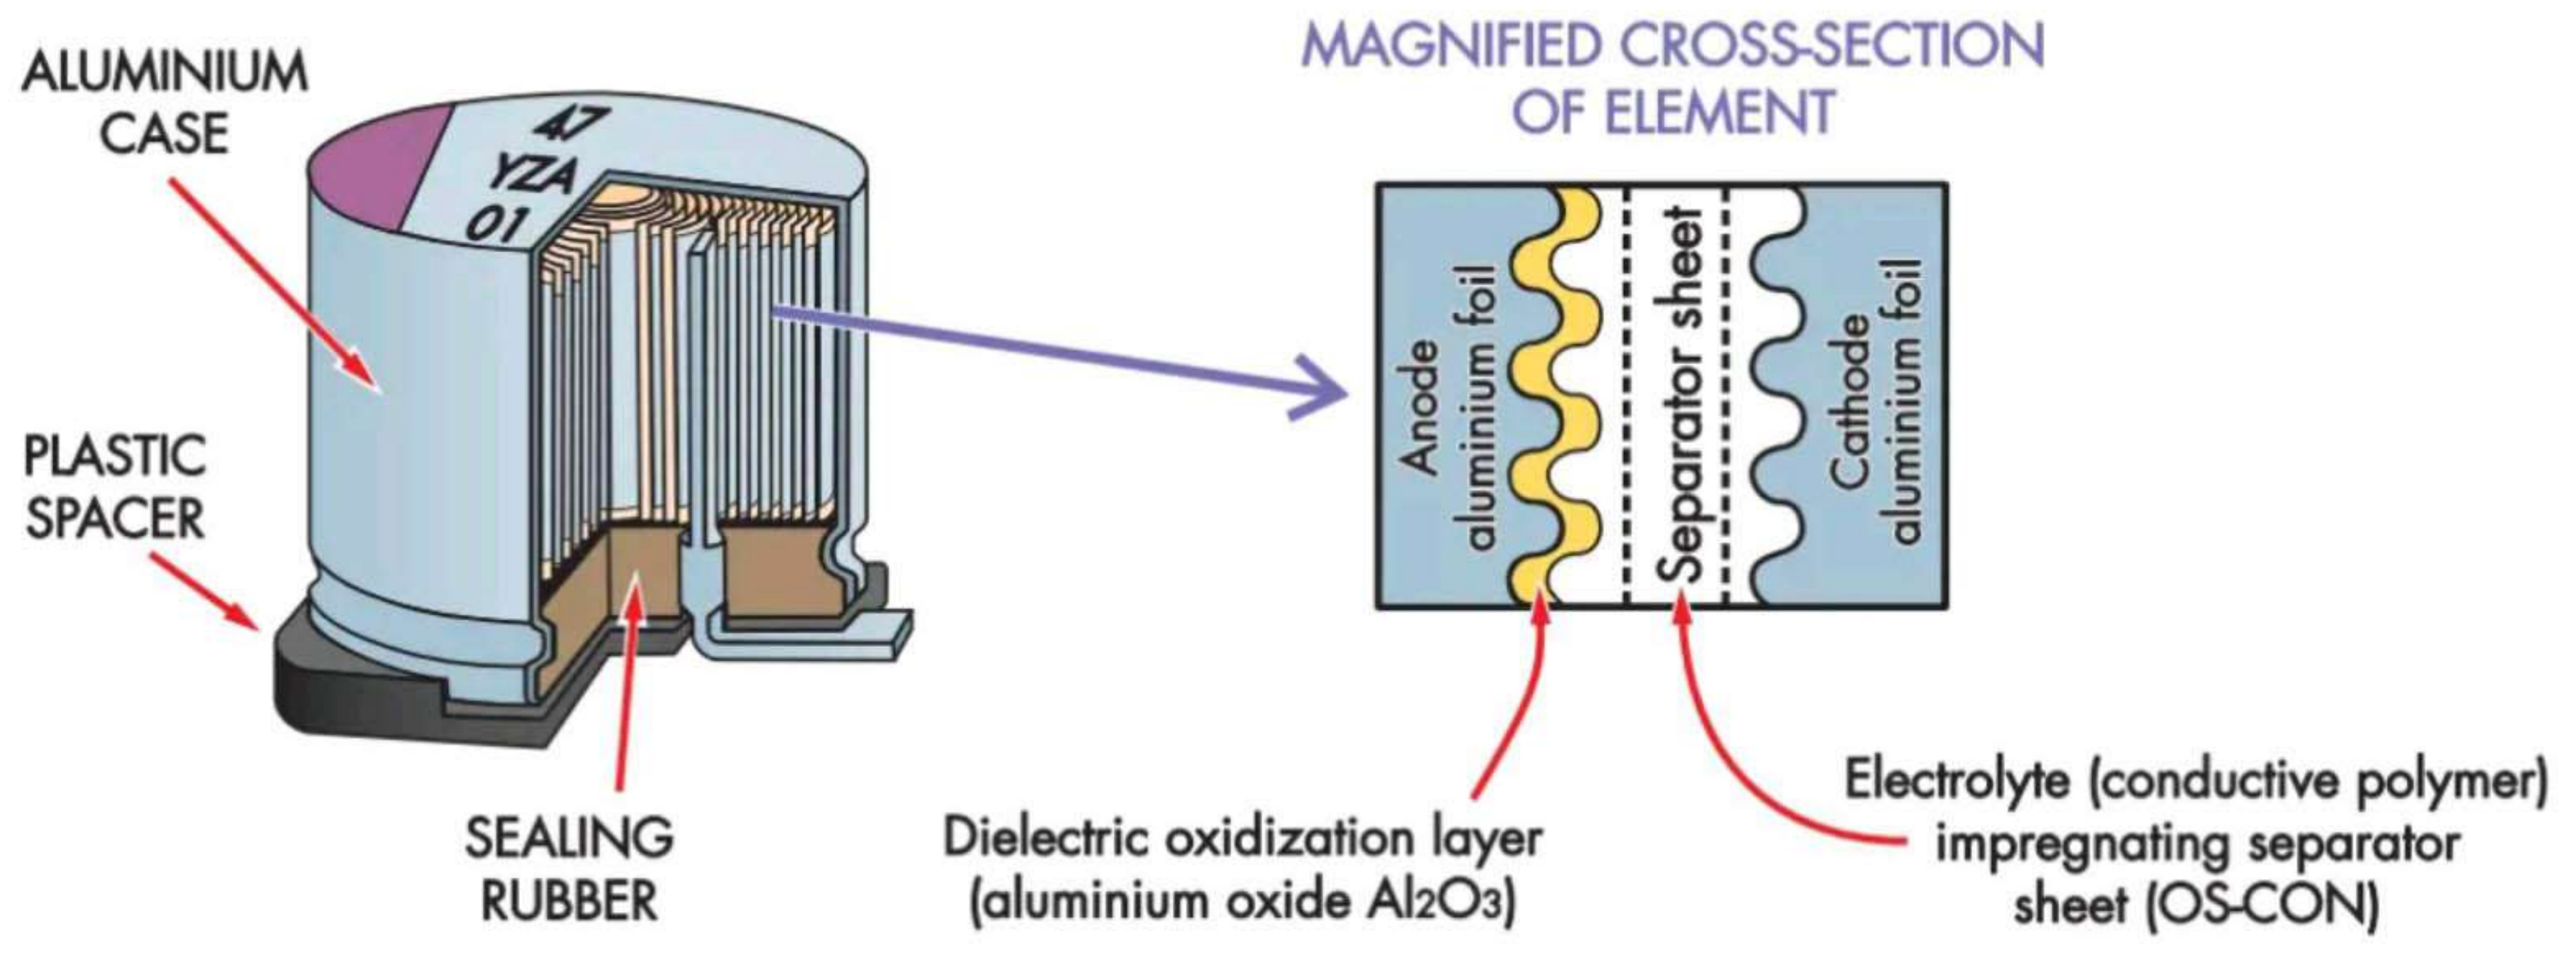
ALUMINIUM CASE
47 YZA 01
PLASTIC SPACER
SEALING RUBBER
MAGNIFIED CROSS-SECTION OF ELEMENT
Anode aluminium foil
Separator sheet
Cathode aluminium foil
Dielectric oxidization layer (aluminium oxide Al2O3)
Electrolyte (conductive polymer) impregnating separator sheet (OS-CON)

Fig.14: SMD polymer electrolytic capacitors are available in various packages including cans like regular electros. The main difference is the use of a separator sheet impregnated with a conductive polymer instead of paper soaked in an electrolyte solution. ***(original source: Panasonic)***

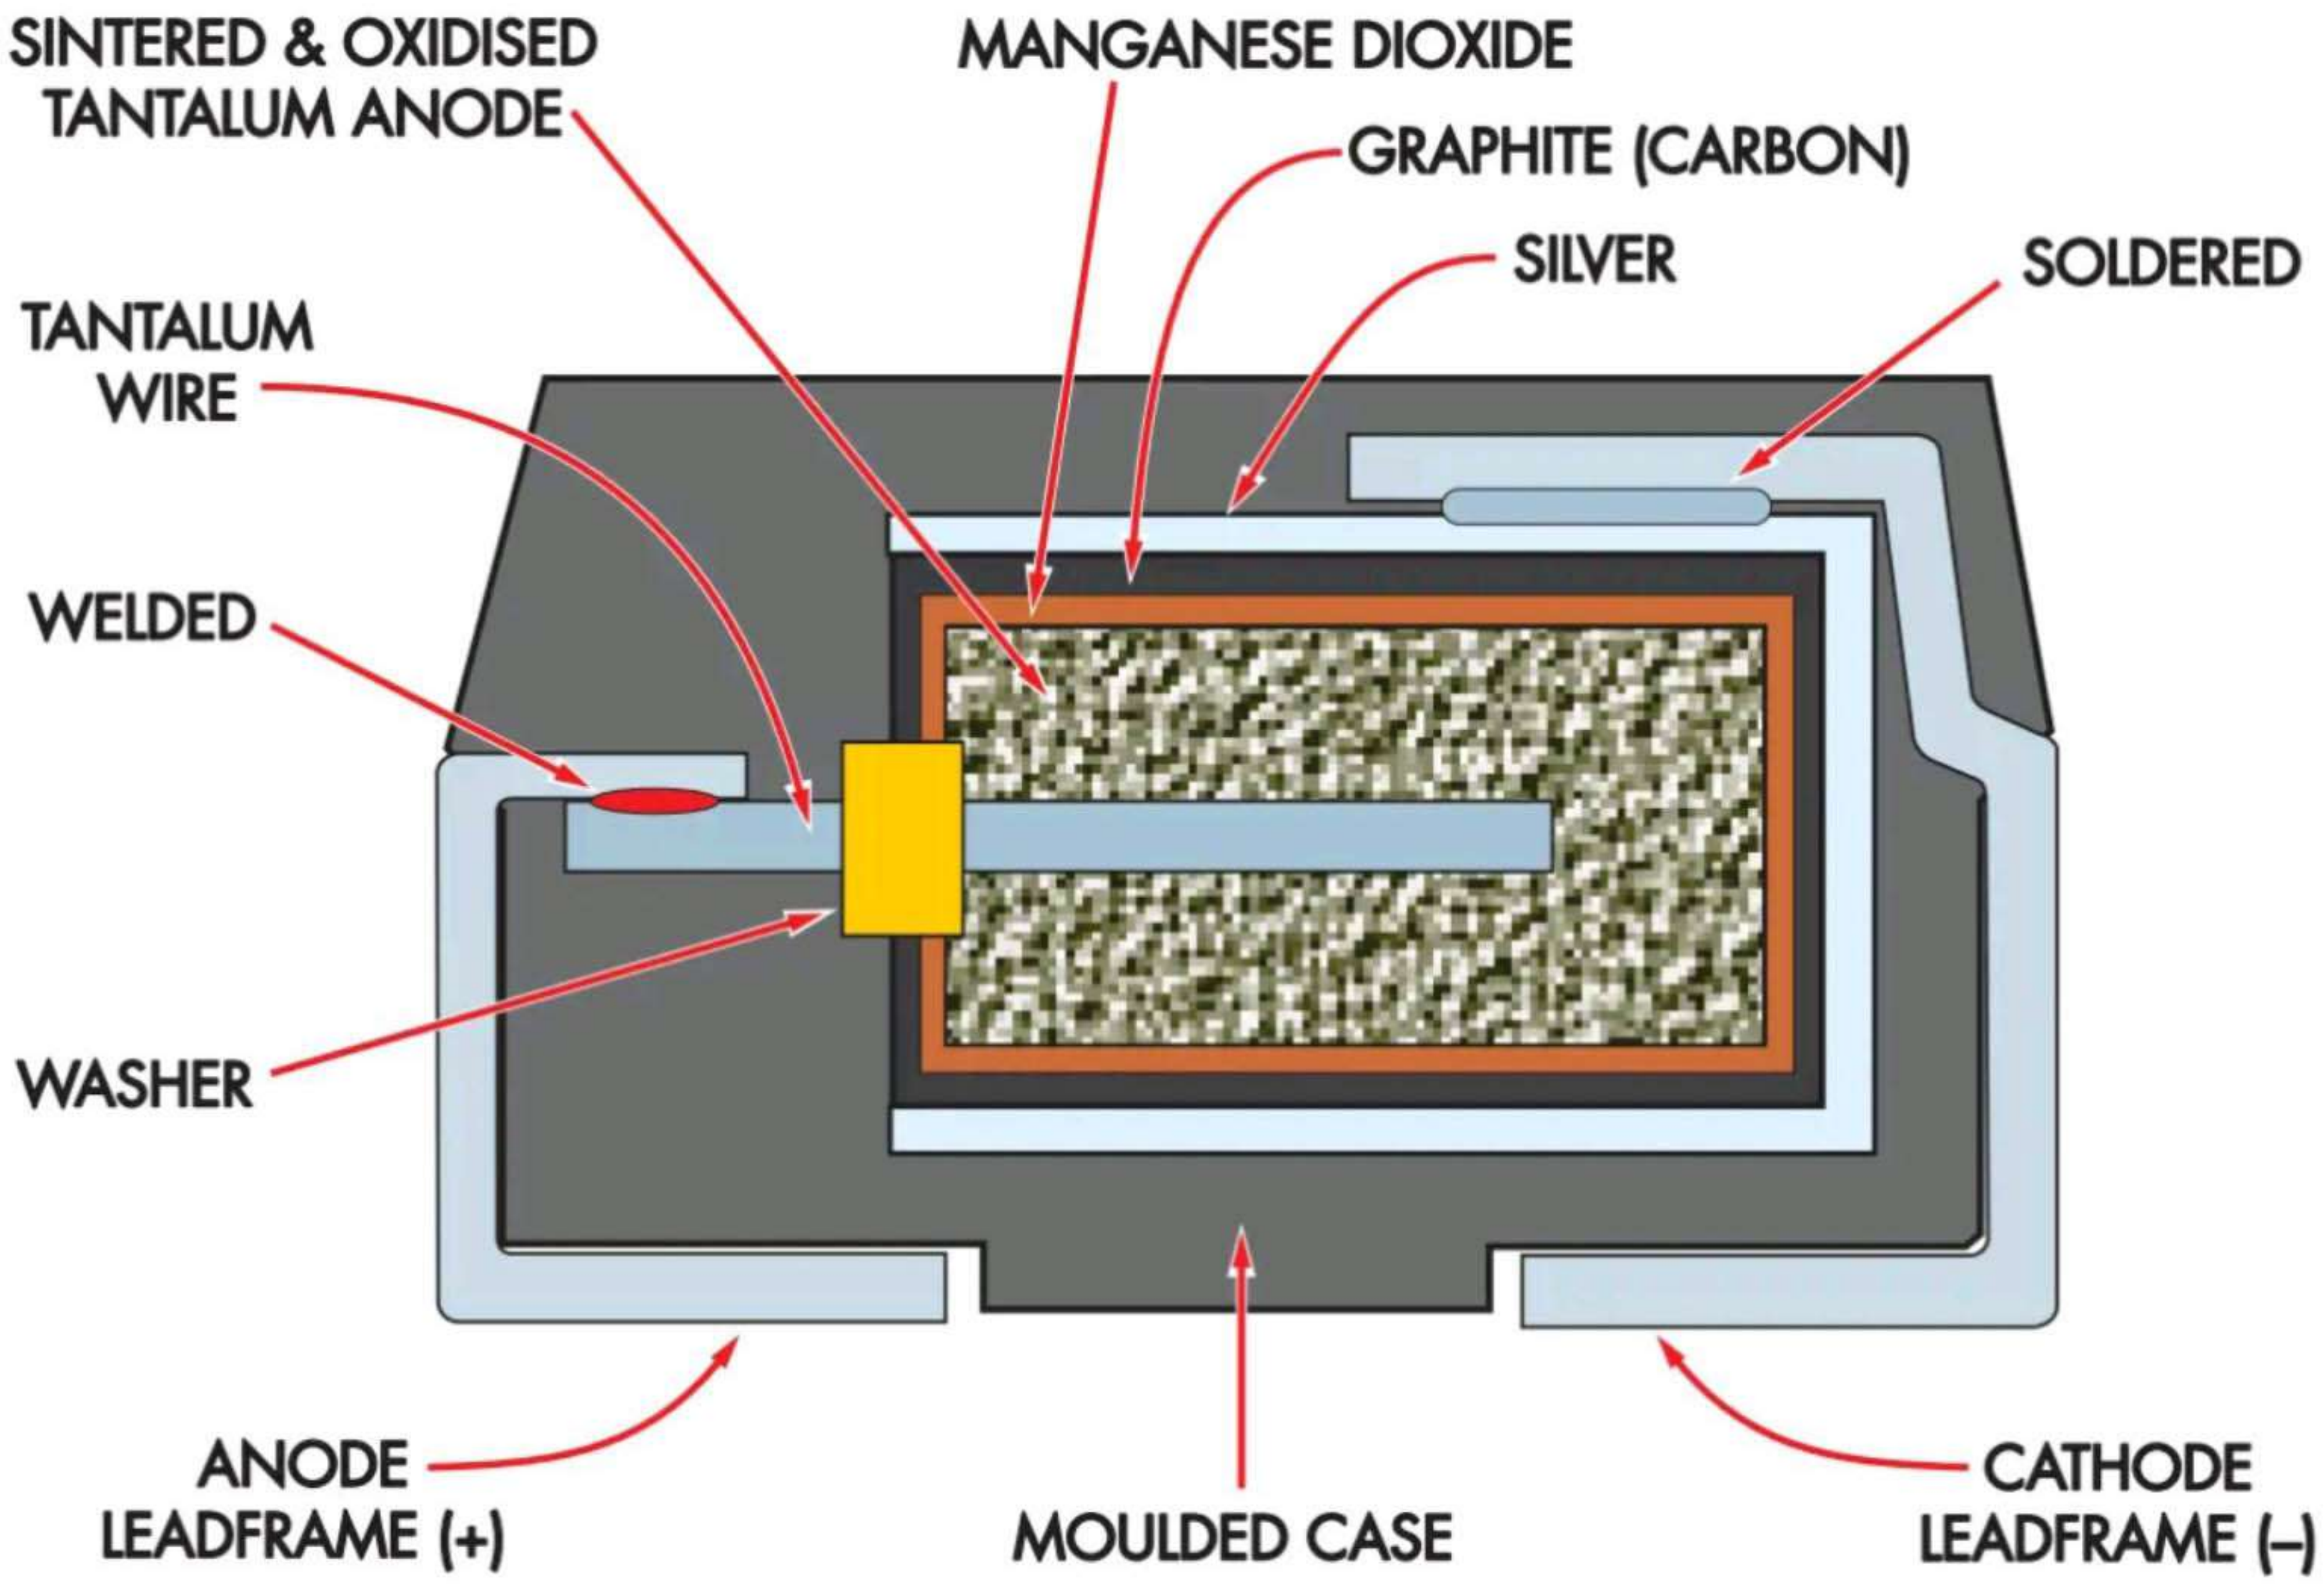
SINTERED & OXIDISED TANTALUM ANODE
MANGANESE DIOXIDE
GRAPHITE (CARBON)
SILVER
SOLDERED
TANTALUM WIRE
WELDED
WASHER
ANODE LEADFRAME (+)
MOULDED CASE
CATHODE LEADFRAME (–)

Fig.15: SMD tantalum electros are made similarly to through-hole types, but the sintered tantalum grains are formed in a rectangular prism shape to create a more compact and convenient package. ***(original source: Wikipedia)***

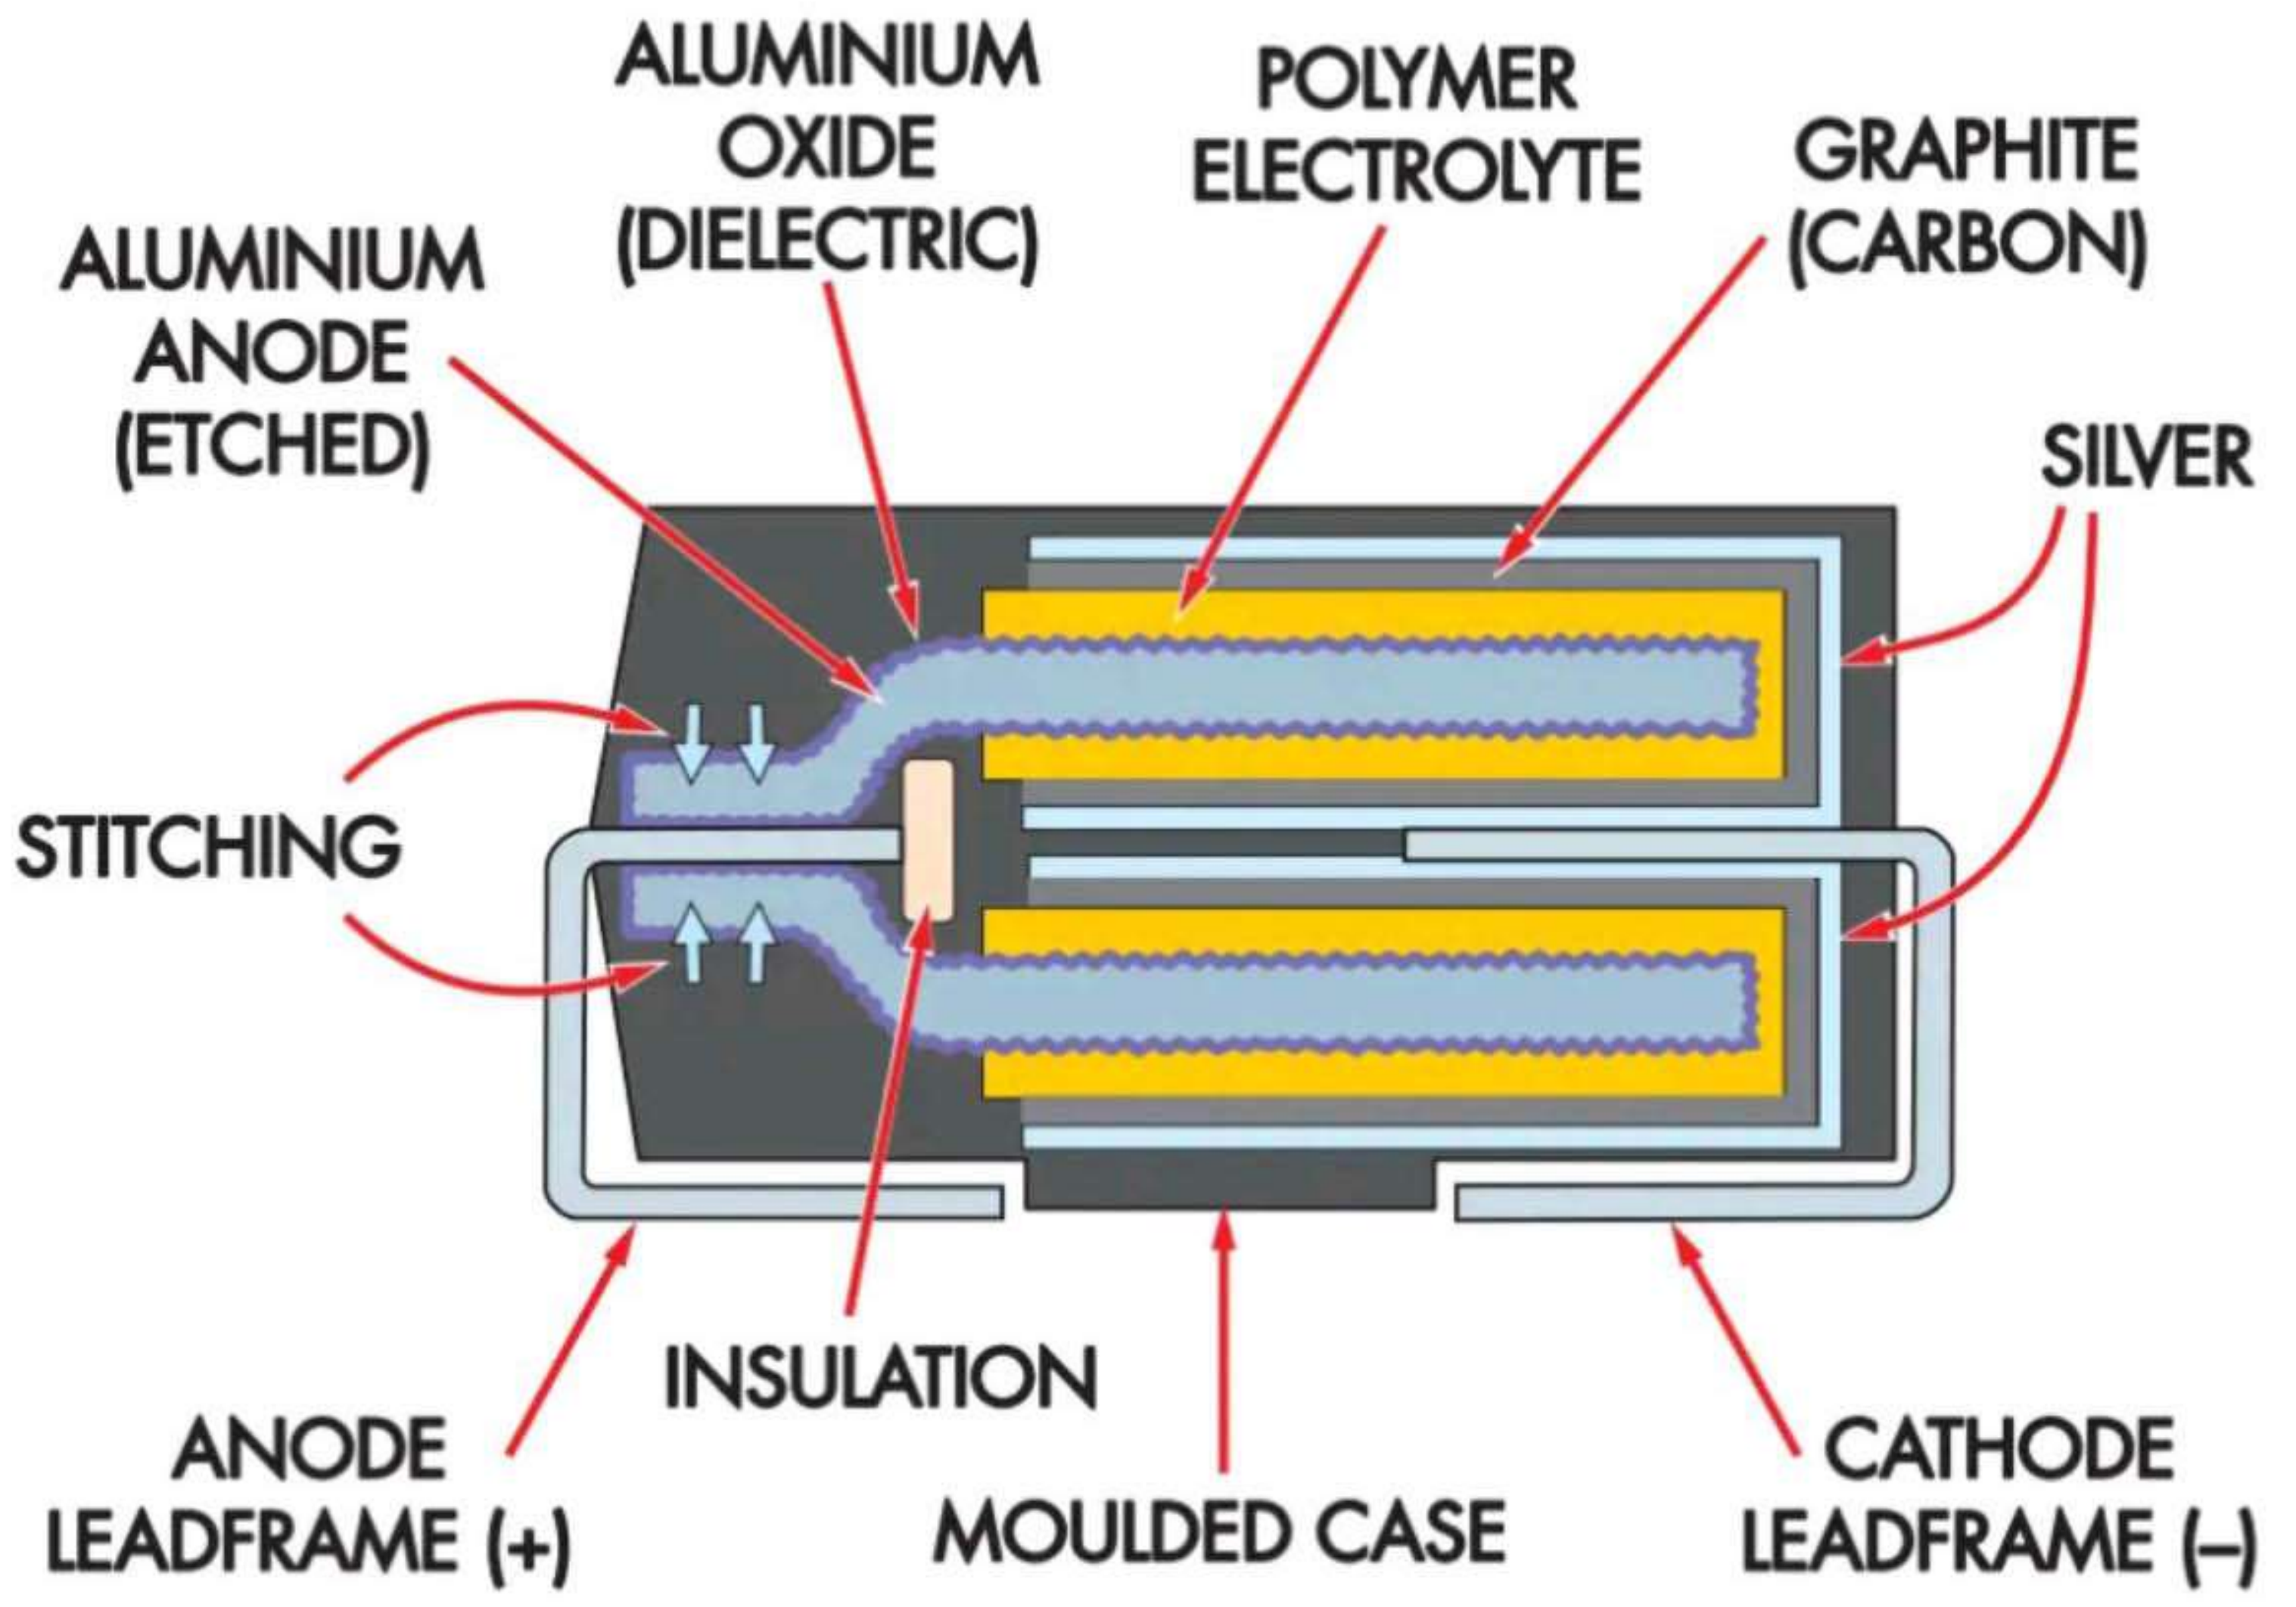
ALUMINIUM ANODE (ETCHED)
ALUMINIUM OXIDE (DIELECTRIC)
POLYMER ELECTROLYTE
GRAPHITE (CARBON)
SILVER
STITCHING
INSULATION
ANODE LEADFRAME (+)
MOULDED CASE
CATHODE LEADFRAME (–)

Fig.16: the same type of semi-rectangular SMD package can also be used to house polymer aluminium electrolytic capacitors. ***(original source: Wikipedia)***

Some through-hole tantalum capacitors ranging from 3.3μF to 47μF. For these capacitors, polarity is indicated by a plus sign (+) on one side of the body.

Current is passed through the capacitor to form the required insulating layer, up to a voltage somewhat above the desired rating (which determines the oxide layer thickness).

The process is slightly different for tantalum and polymer capacitors; SMD tantalum and polymer capacitors in rectangular prism packages may be made in layers rather than wound. But the result is much the same: a metal conductor with a large surface area separated from a conductive electrolyte only by a very thin oxide layer (see Figs.14-16).

If the leads were only connected to the conductive foils at one point each, the ESR and ESL of the capacitor would be poor, as current must flow along a spiral path to reach the inner and outer layers of the capacitor.

For this reason, all but the most basic electros usually have extra conductive paths giving current a 'short cut' to move between the layers of the capacitor.

Higher performance electros also have the tabs 'stitched' into the foil at multiple points to reduce resistance and improve conductivity between them, in addition to having numerous tabs spread throughout the roll, that all join to one of the two leads.

Plastic film capacitors

Plastic film capacitors are not used as much commercially these days since ceramic capacitors are much cheaper and are available with very low ESR and ESL figures. However, in cases where linearity or safety are essential, plastic films are still widely used.

Most plastic film capacitors have better linearity than all but the best (NP0/C0G) ceramics or mica capacitors, and they can be designed to fail gracefully (going open-circuit).

Plastic film capacitors are available from a few dozen picofarads up to a few tens of microfarads, and have voltage ratings ranging from around 16V up to several thousand volts.

The failure mode is vital in mains applications, where capacitors are connected between Active and Neutral or Active and Earth. If they were to fail short circuit, a fire could result, or it could be a shock hazard. While ceramic X/Y-class capacitors exist, generally, higher-value mains-rated (X/Y) capacitors use either polyester (PET), polypropylene (PP) or polycarbonate (PC) films.

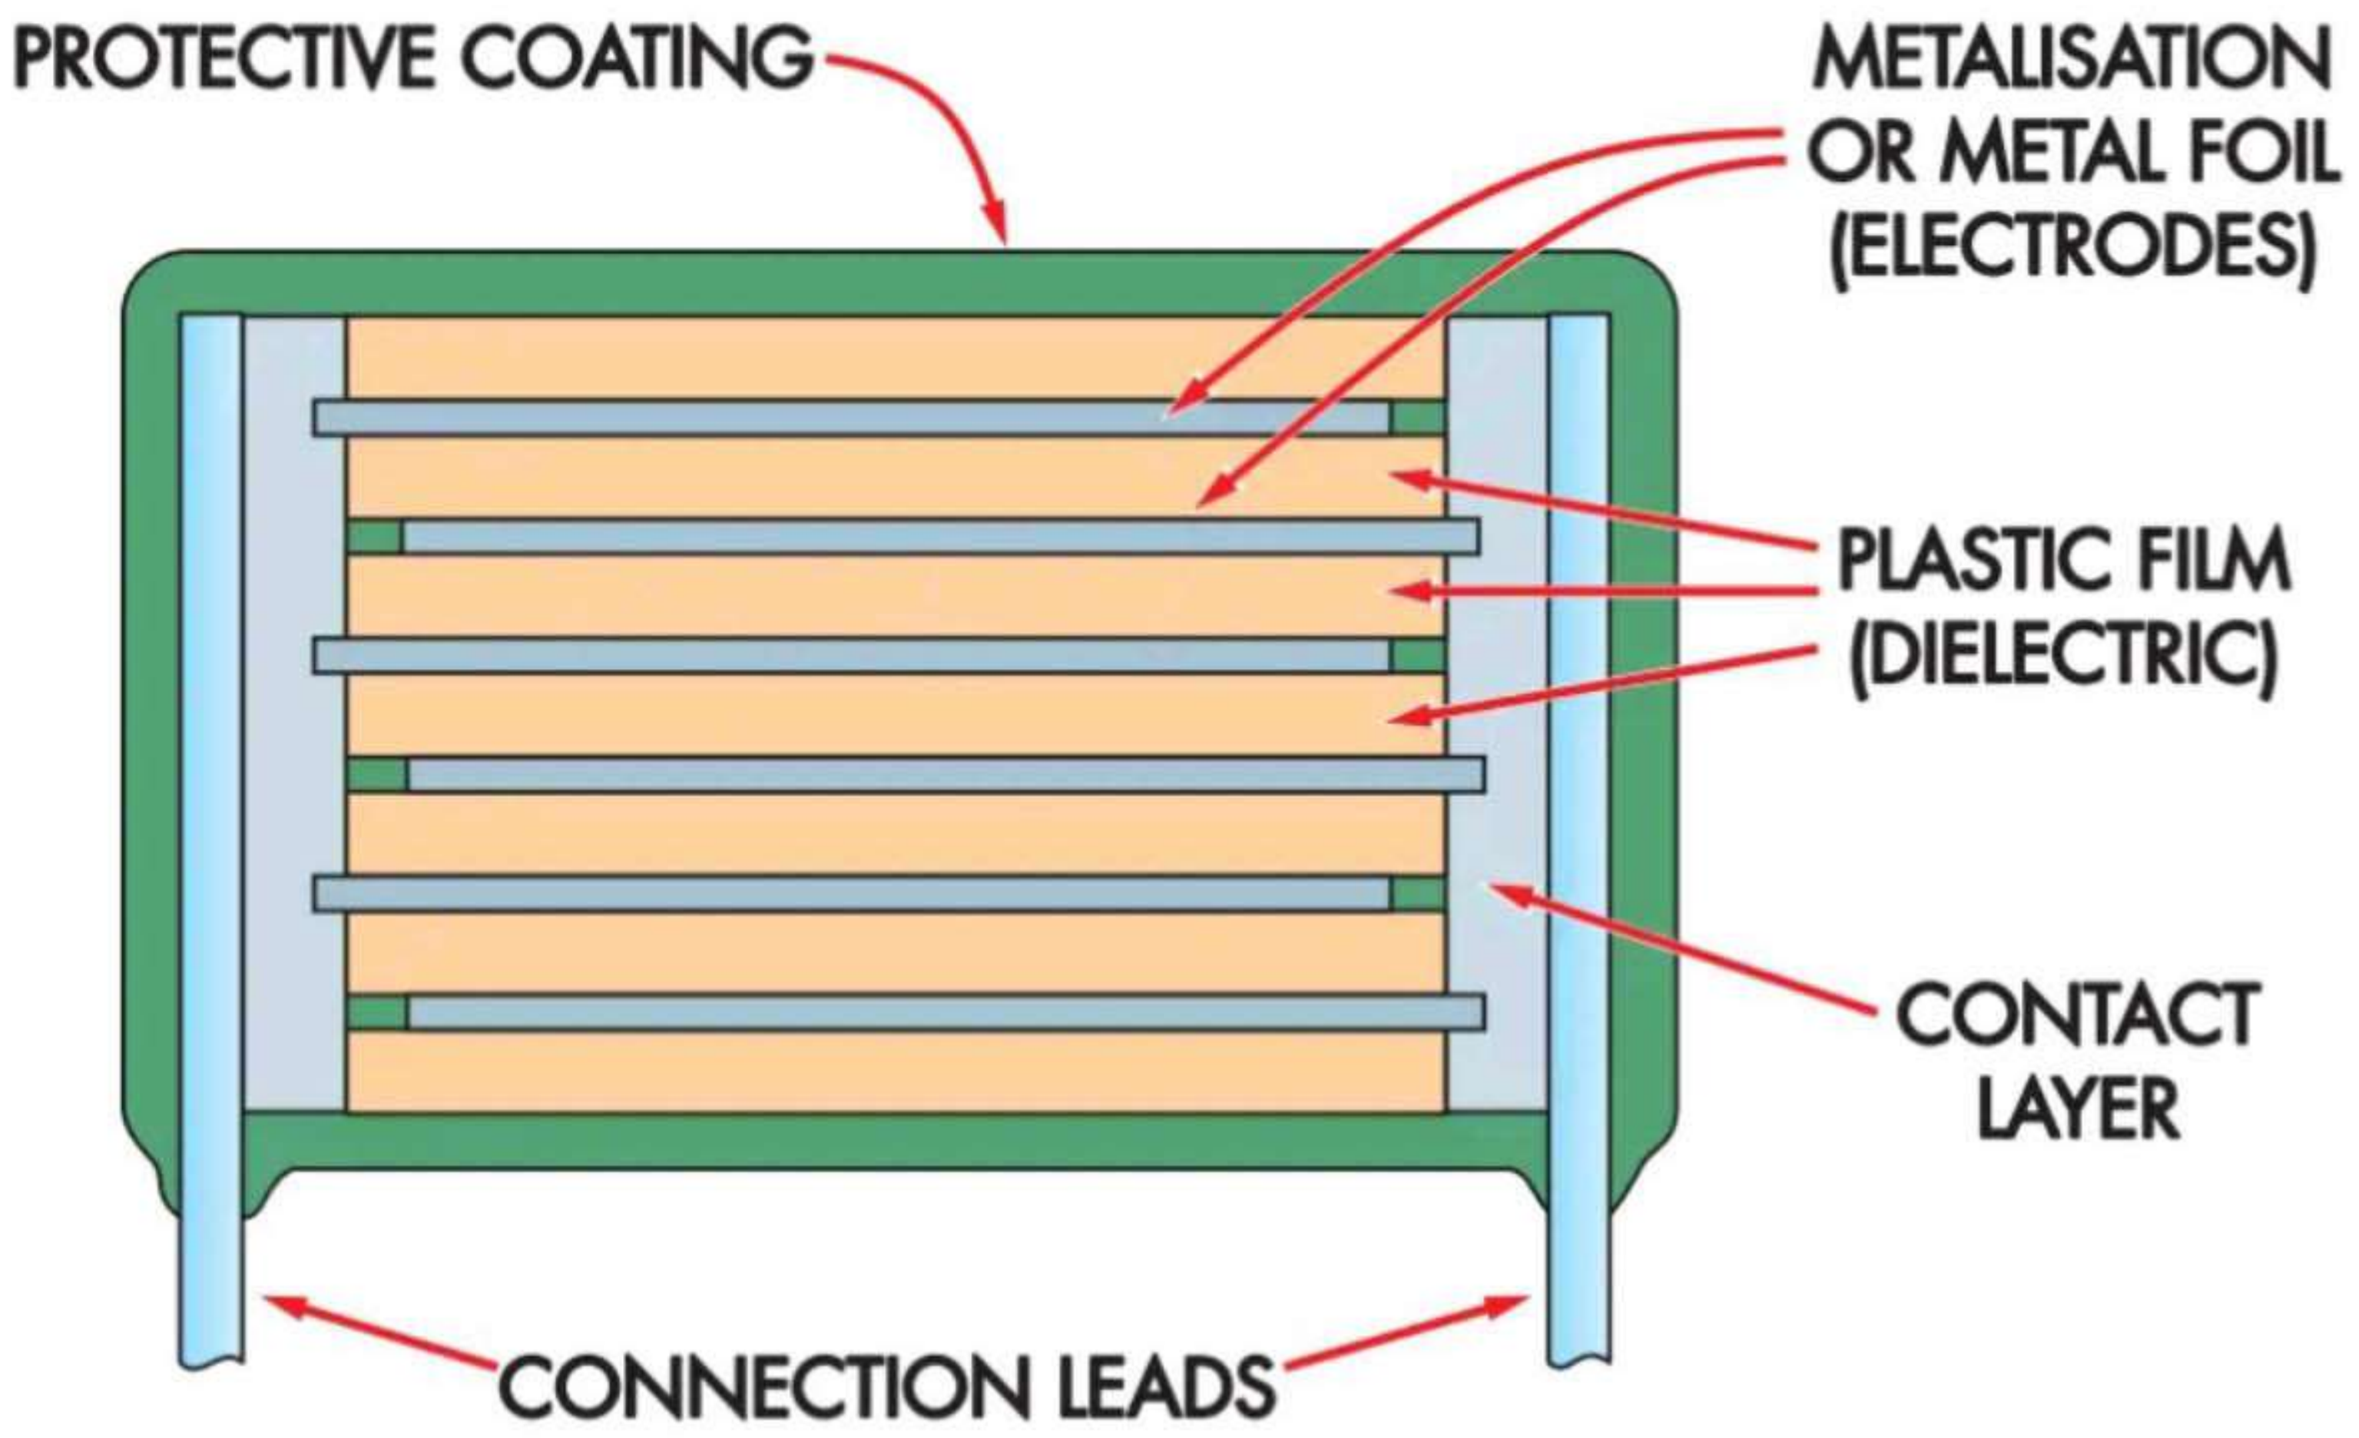
PROTECTIVE COATING
METALISATION OR METAL FOIL (ELECTRODES)
PLASTIC FILM (DIELECTRIC)
CONTACT LAYER
CONNECTION LEADS

Fig.17: plastic capacitor construction is similar to ceramic, with alternating layers of plastic (the dielectric) and conductive metal film or foil in between, staggered to create many capacitors in parallel (original source: Wikipedia).

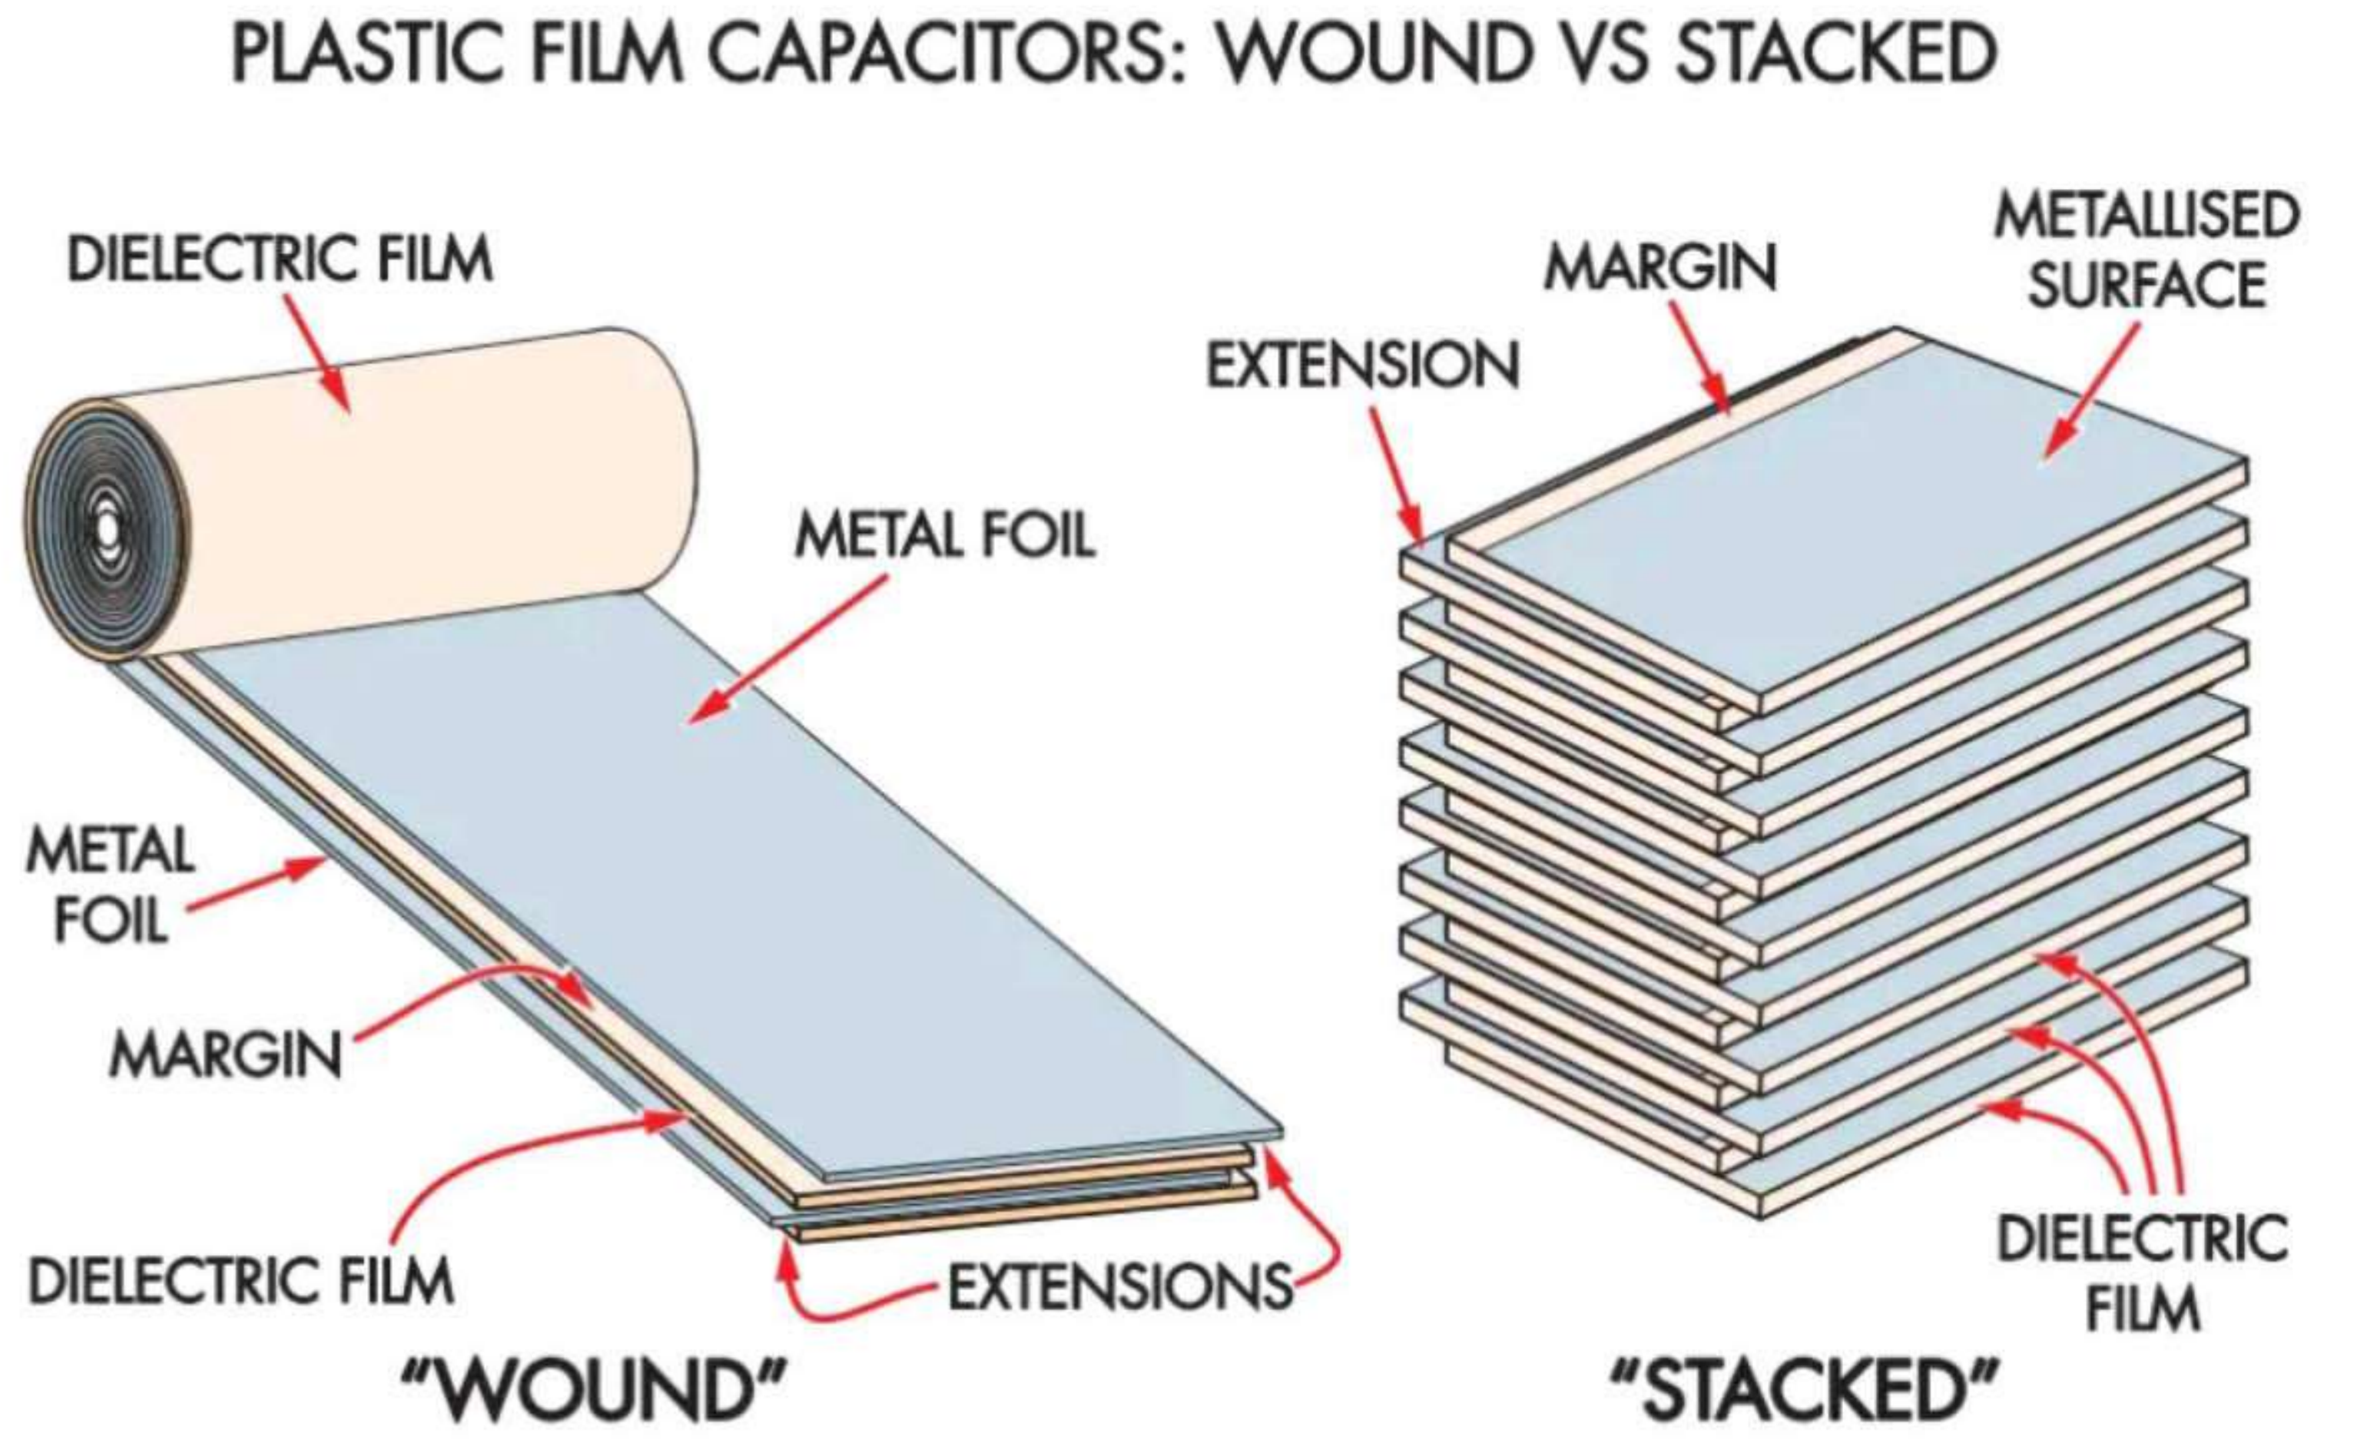
PLASTIC FILM CAPACITORS: WOUND VS STACKED
DIELECTRIC FILM
METAL FOIL
METAL FOIL
MARGIN
DIELECTRIC FILM
EXTENSIONS
"WOUND"
MARGIN
METALLISED SURFACE
EXTENSION
DIELECTRIC FILM
"STACKED"

Fig.18: plastic capacitors can be made from stacks of sheets, similarly to ceramics, or from rolled-up strips, similarly to electrolytics. The stacking process tends to be more time-consuming and expensive, but it can give better density, so it is used for some SMD plastic capacitors.

Like ceramic capacitors, the plastic dielectric used has a significant effect on capacitor properties. And similarly, the plastics with the lowest dielectric constants that result in the bulkiest capacitors (like polypropylene and polystyrene [PS]), tend to have the best performance figures, such as good linearity factors and low dissipation factors.

Other plastics used for capacitors include polyphenylene sulfide (PPS), polyethylene naphthalate (PEN) and polytetrafluoroethylene (PTFE).

One interesting property of metallised plastic film capacitors is 'self-healing'. This is where a physical defect or the application of excessive voltage might damage the capacitor, but it will not fail entirely; instead, a small area of the metallisation burns away, reducing the capacitance by (hopefully) a tiny amount – usually not enough to affect its function.

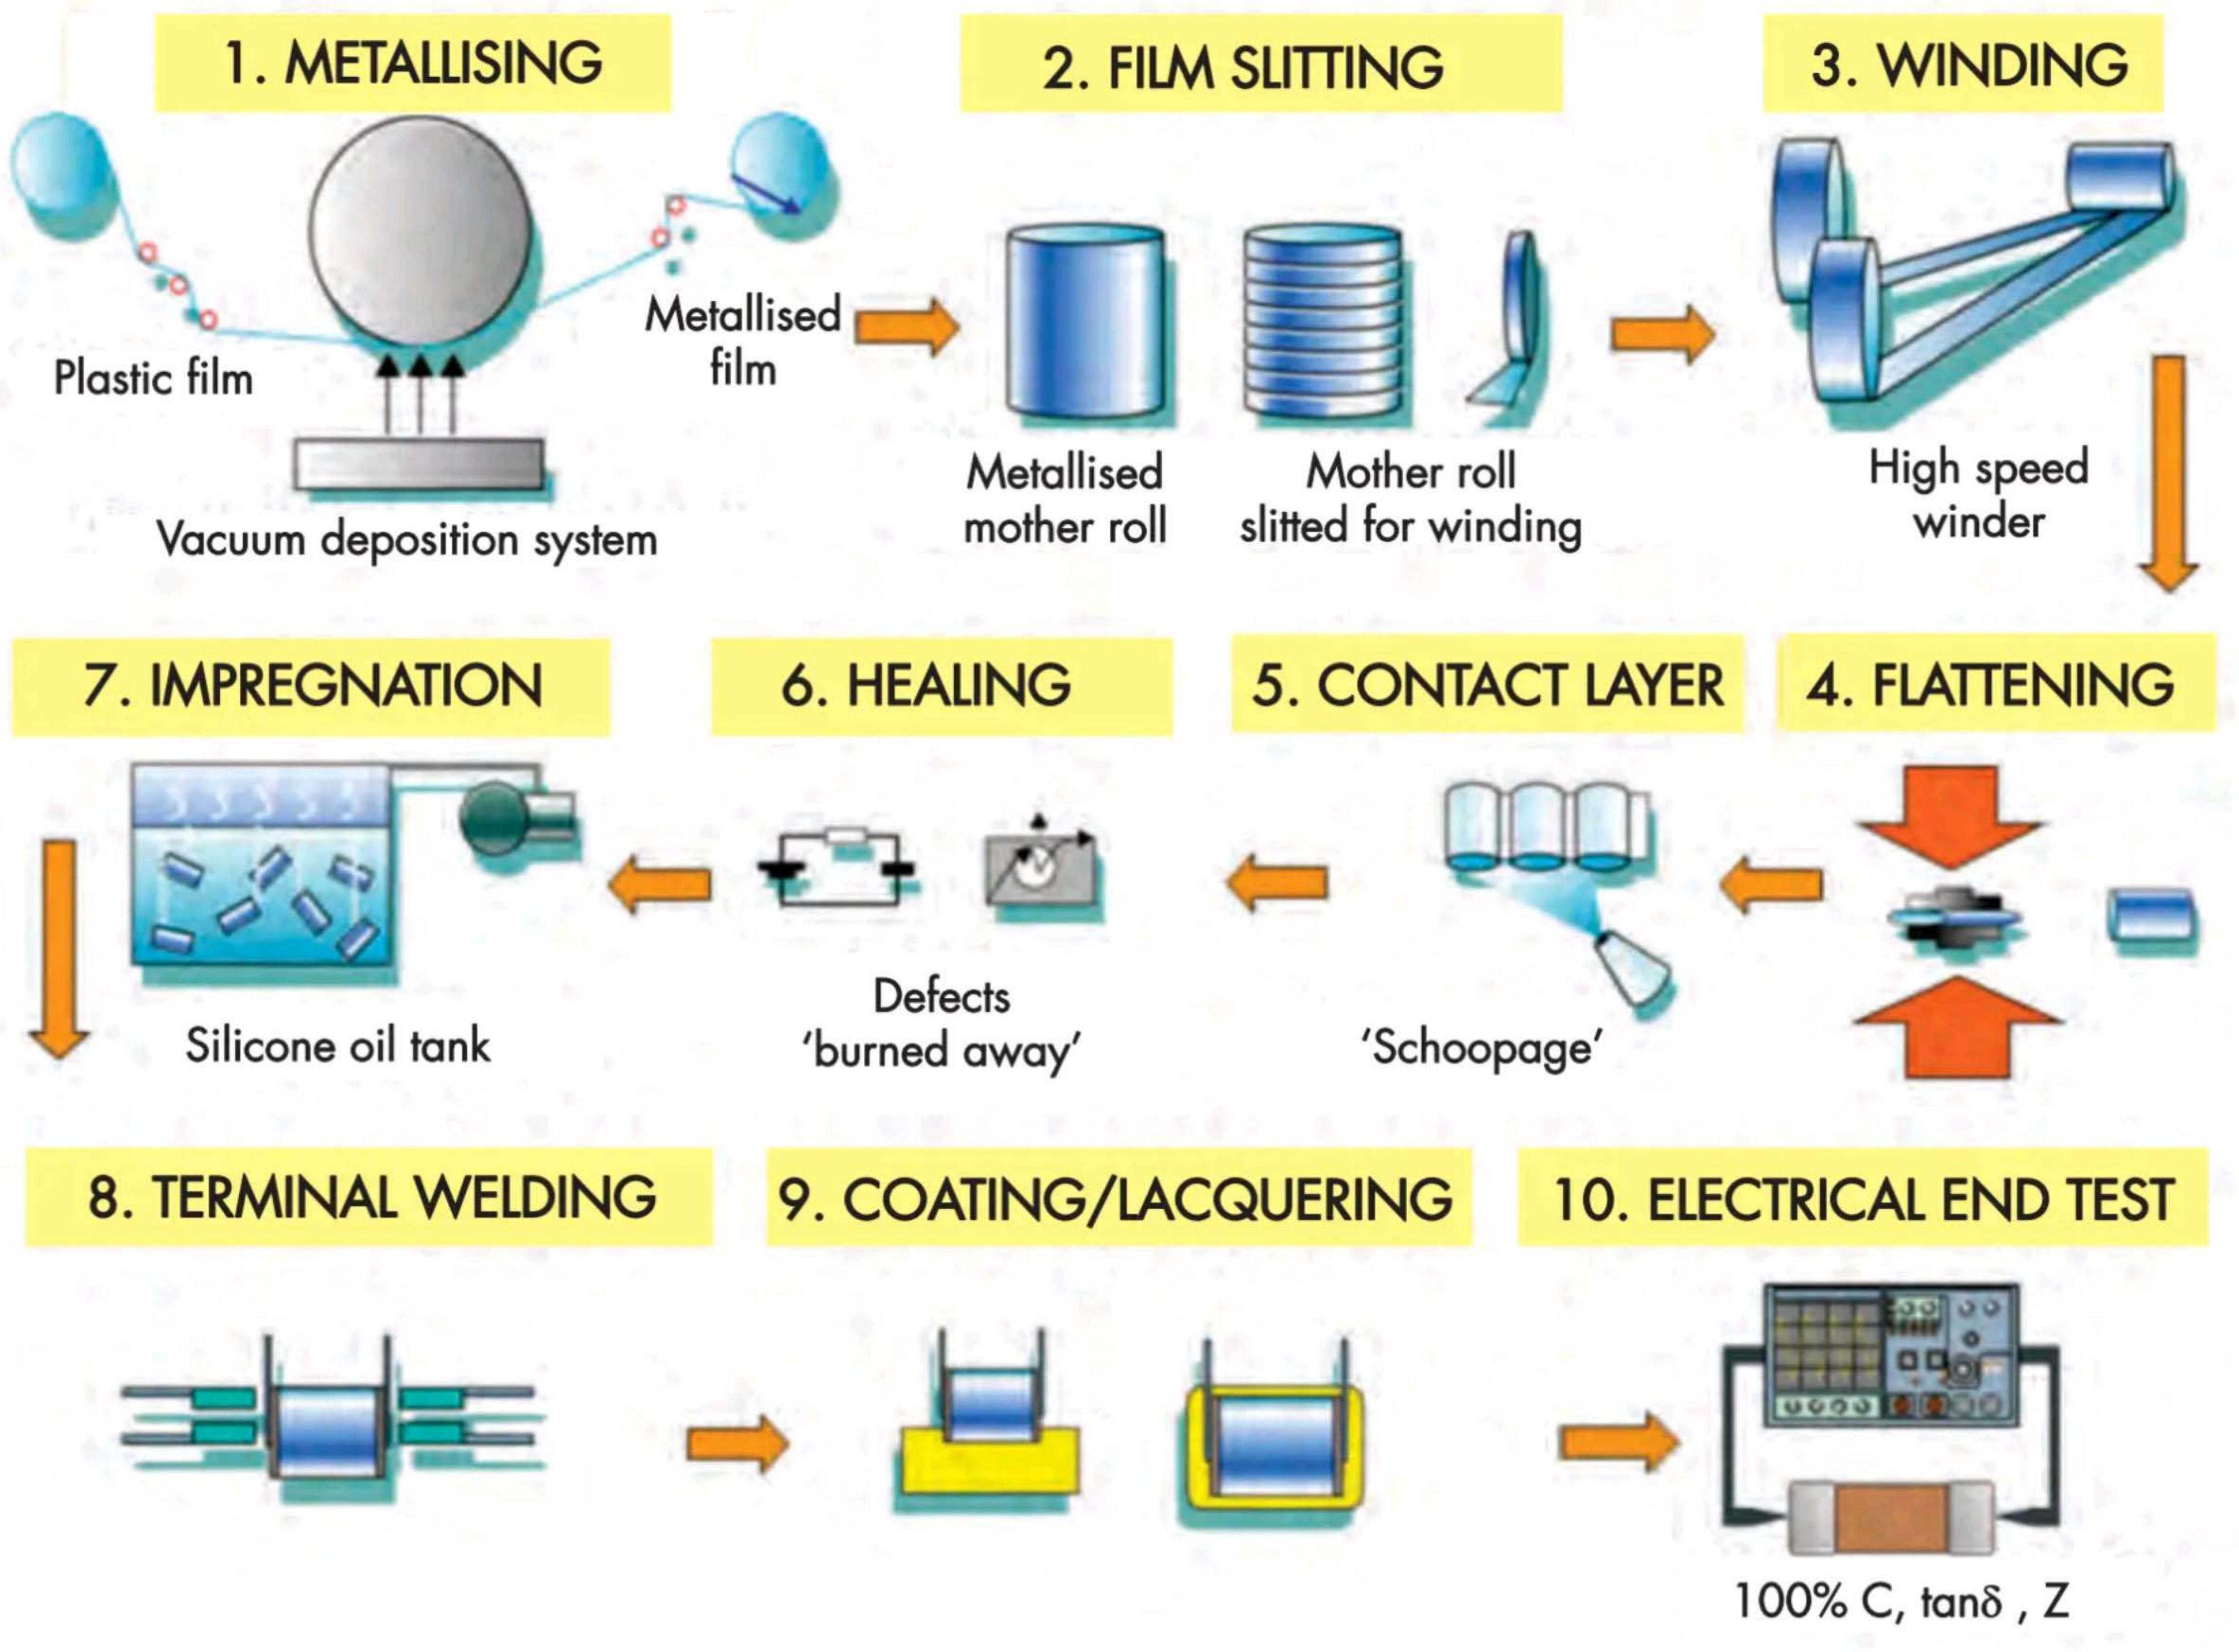
1. METALLISING
Plastic film
Vacuum deposition system
Metallised film
2. FILM SLITTING
Metallised mother roll
Mother roll slitted for winding
3. WINDING
High speed winder
4. FLATTENING
5. CONTACT LAYER
'Schoopage'
6. HEALING
Defects 'burned away'
7. IMPREGNATION
Silicone oil tank
8. TERMINAL WELDING
9. COATING/LACQUERING
10. ELECTRICAL END TEST
100% C, tanδ , Z

Fig.19: the roll manufacturing process for metallised plastic capacitors. While metal foil can be used instead of metallisation, it tends to result in a bulkier capacitor with inferior properties. *(Source: Wikimedia user Elcap)*

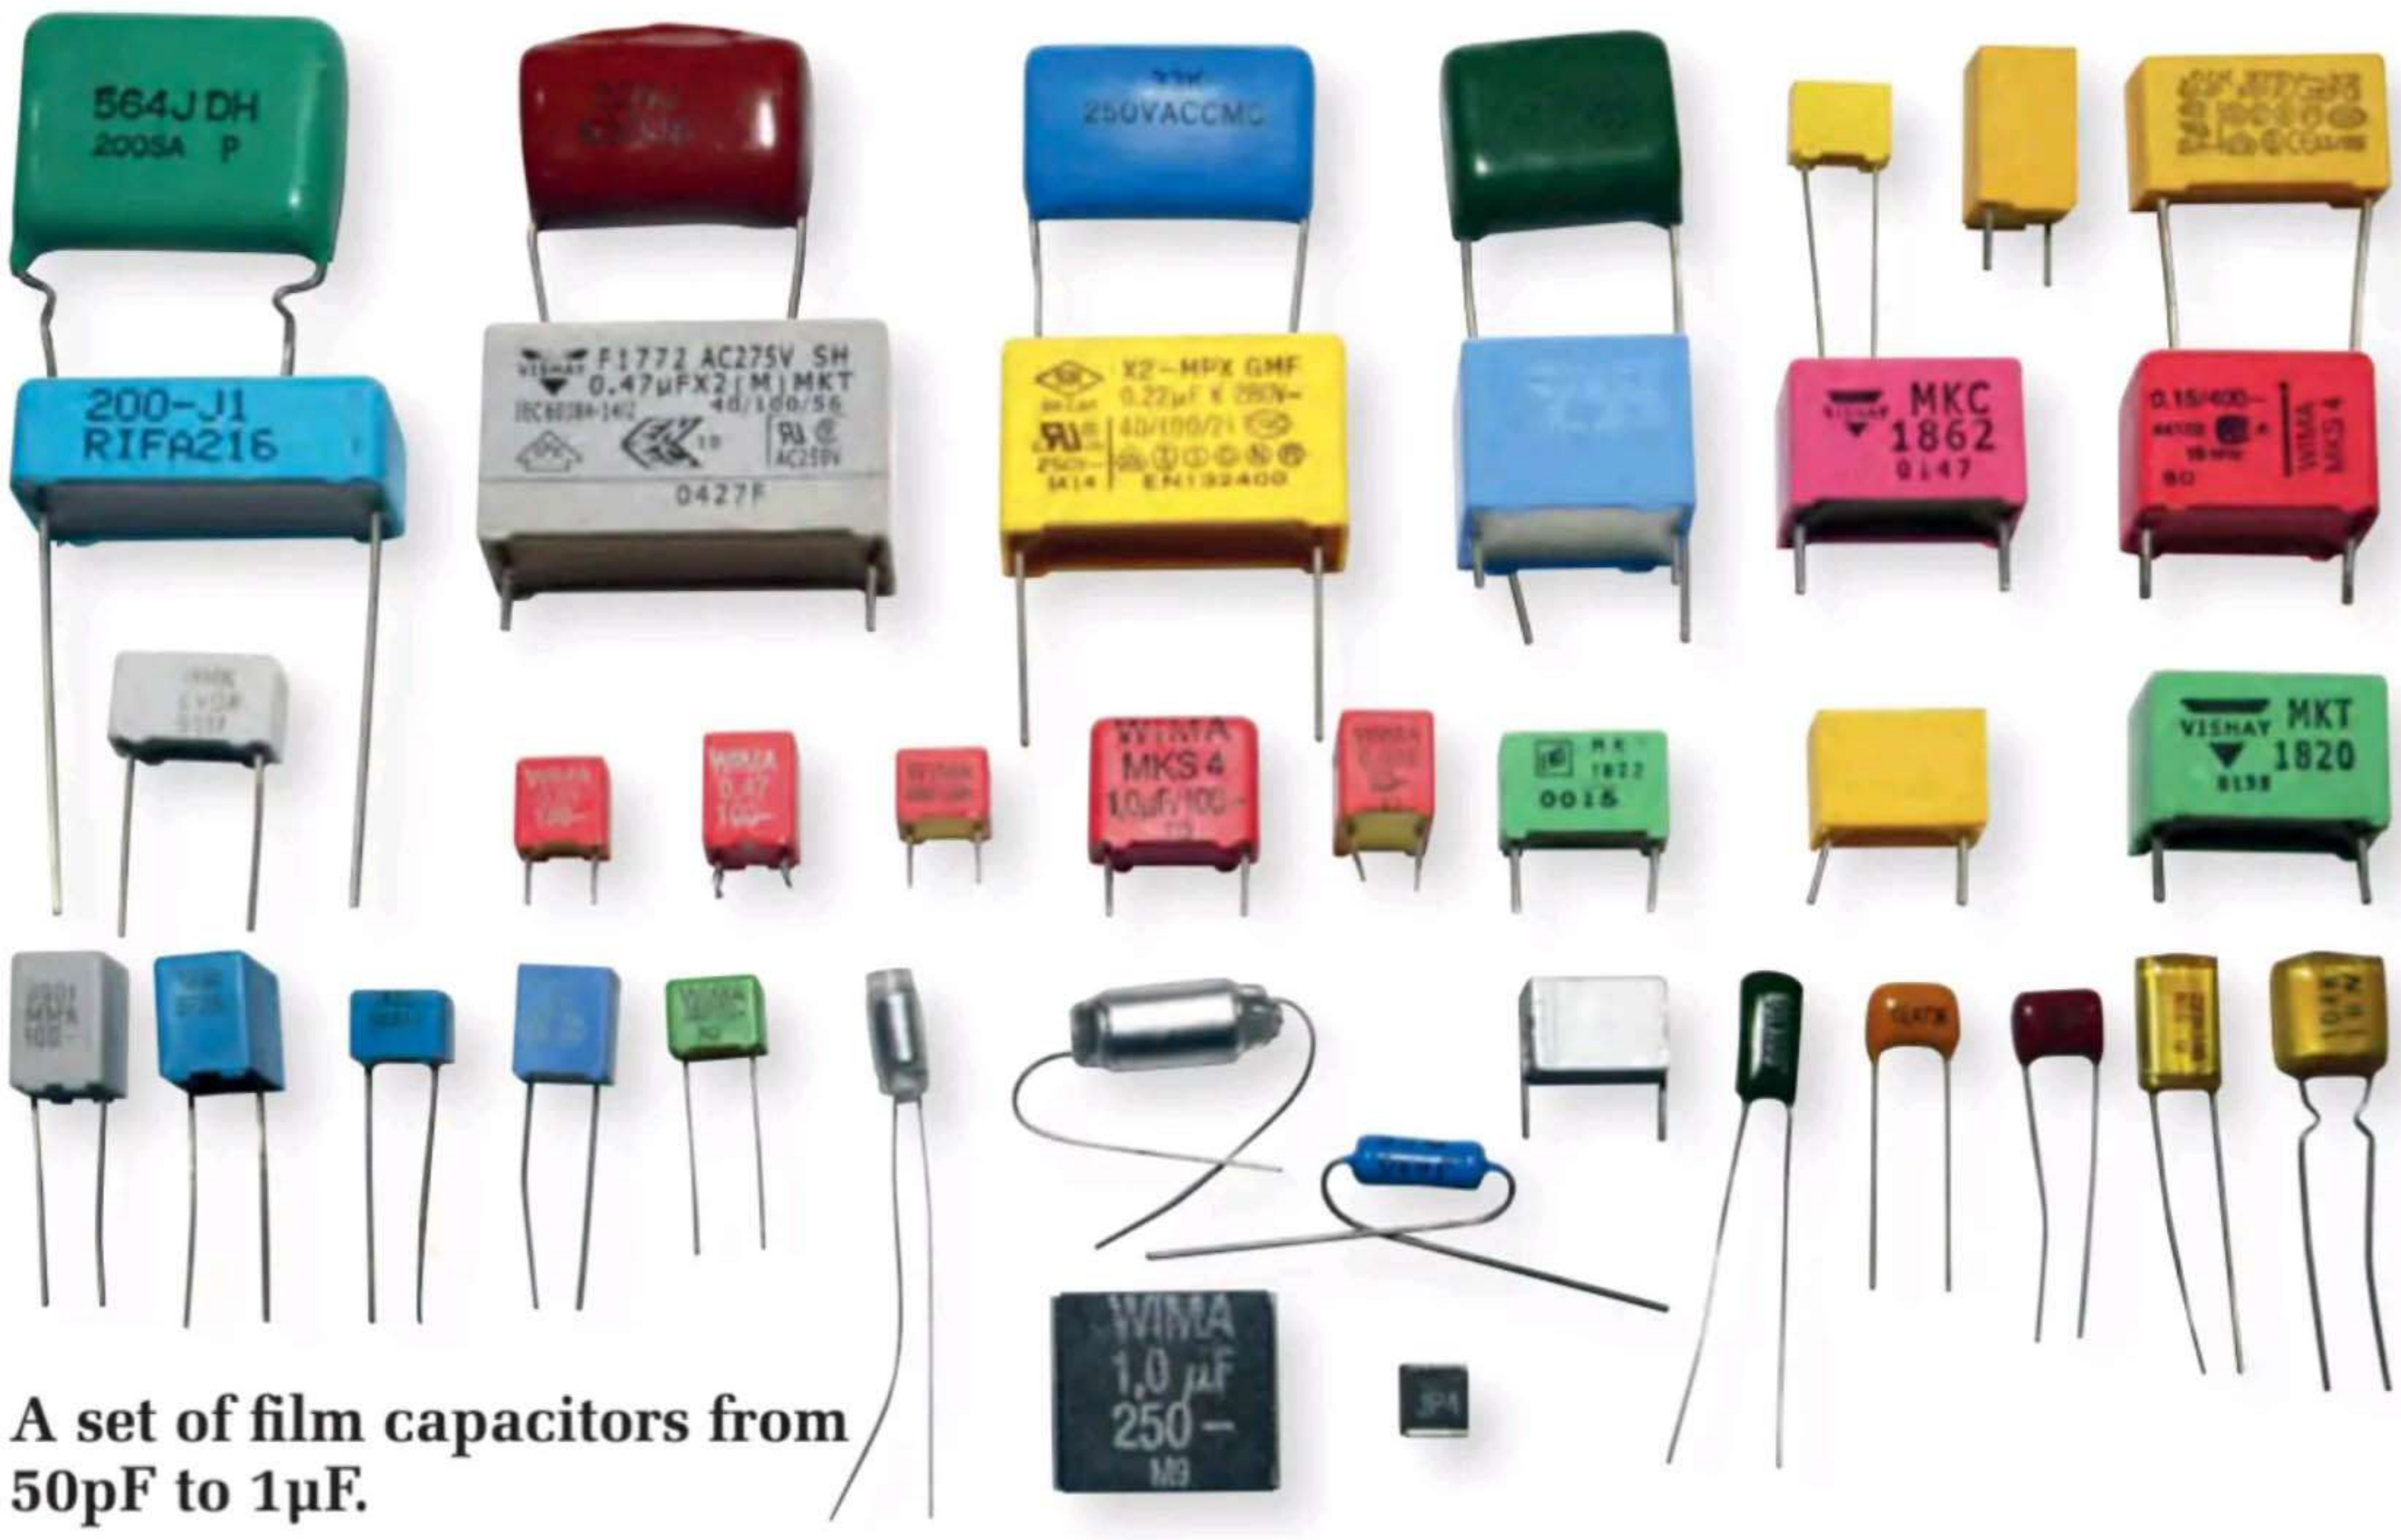

A set of film capacitors from 50pF to 1µF.

This effect is taken advantage of to provide the higher safety margin required for mains-rated capacitors.

Plastic capacitor construction

Like ceramic capacitors, plastic film capacitors usually need many layers to give a usable capacitance value. However, they are not normally formed by deposition methods. Therefore, they must be either stacked or wound (see Figs.17-19).

Stacked capacitors are made in the way you might guess: with alternating layers of conducting foils and dielectric films. The conducting foils are staggered so that when they are joined along each edge, they form interlocking 'combs' and thus effectively, many single-layer capacitors in parallel. Making capacitors that way is time-consuming and expensive, though.

Rather than using metal foils, the dielectric can also be coated with a film of conductive material which provides a thinner and more uniform layer, improving performance and allowing more layers to be packed into the same space for higher density.

Wound plastic film capacitors are made similarly to electrolytics as described above. A sandwich is created with the dielectric film between two strips of metal foil, with the strips slightly offset. This assembly is wound up into a roll, and in most cases, the roll is squashed flat to better fit into a rectangular prism shape.

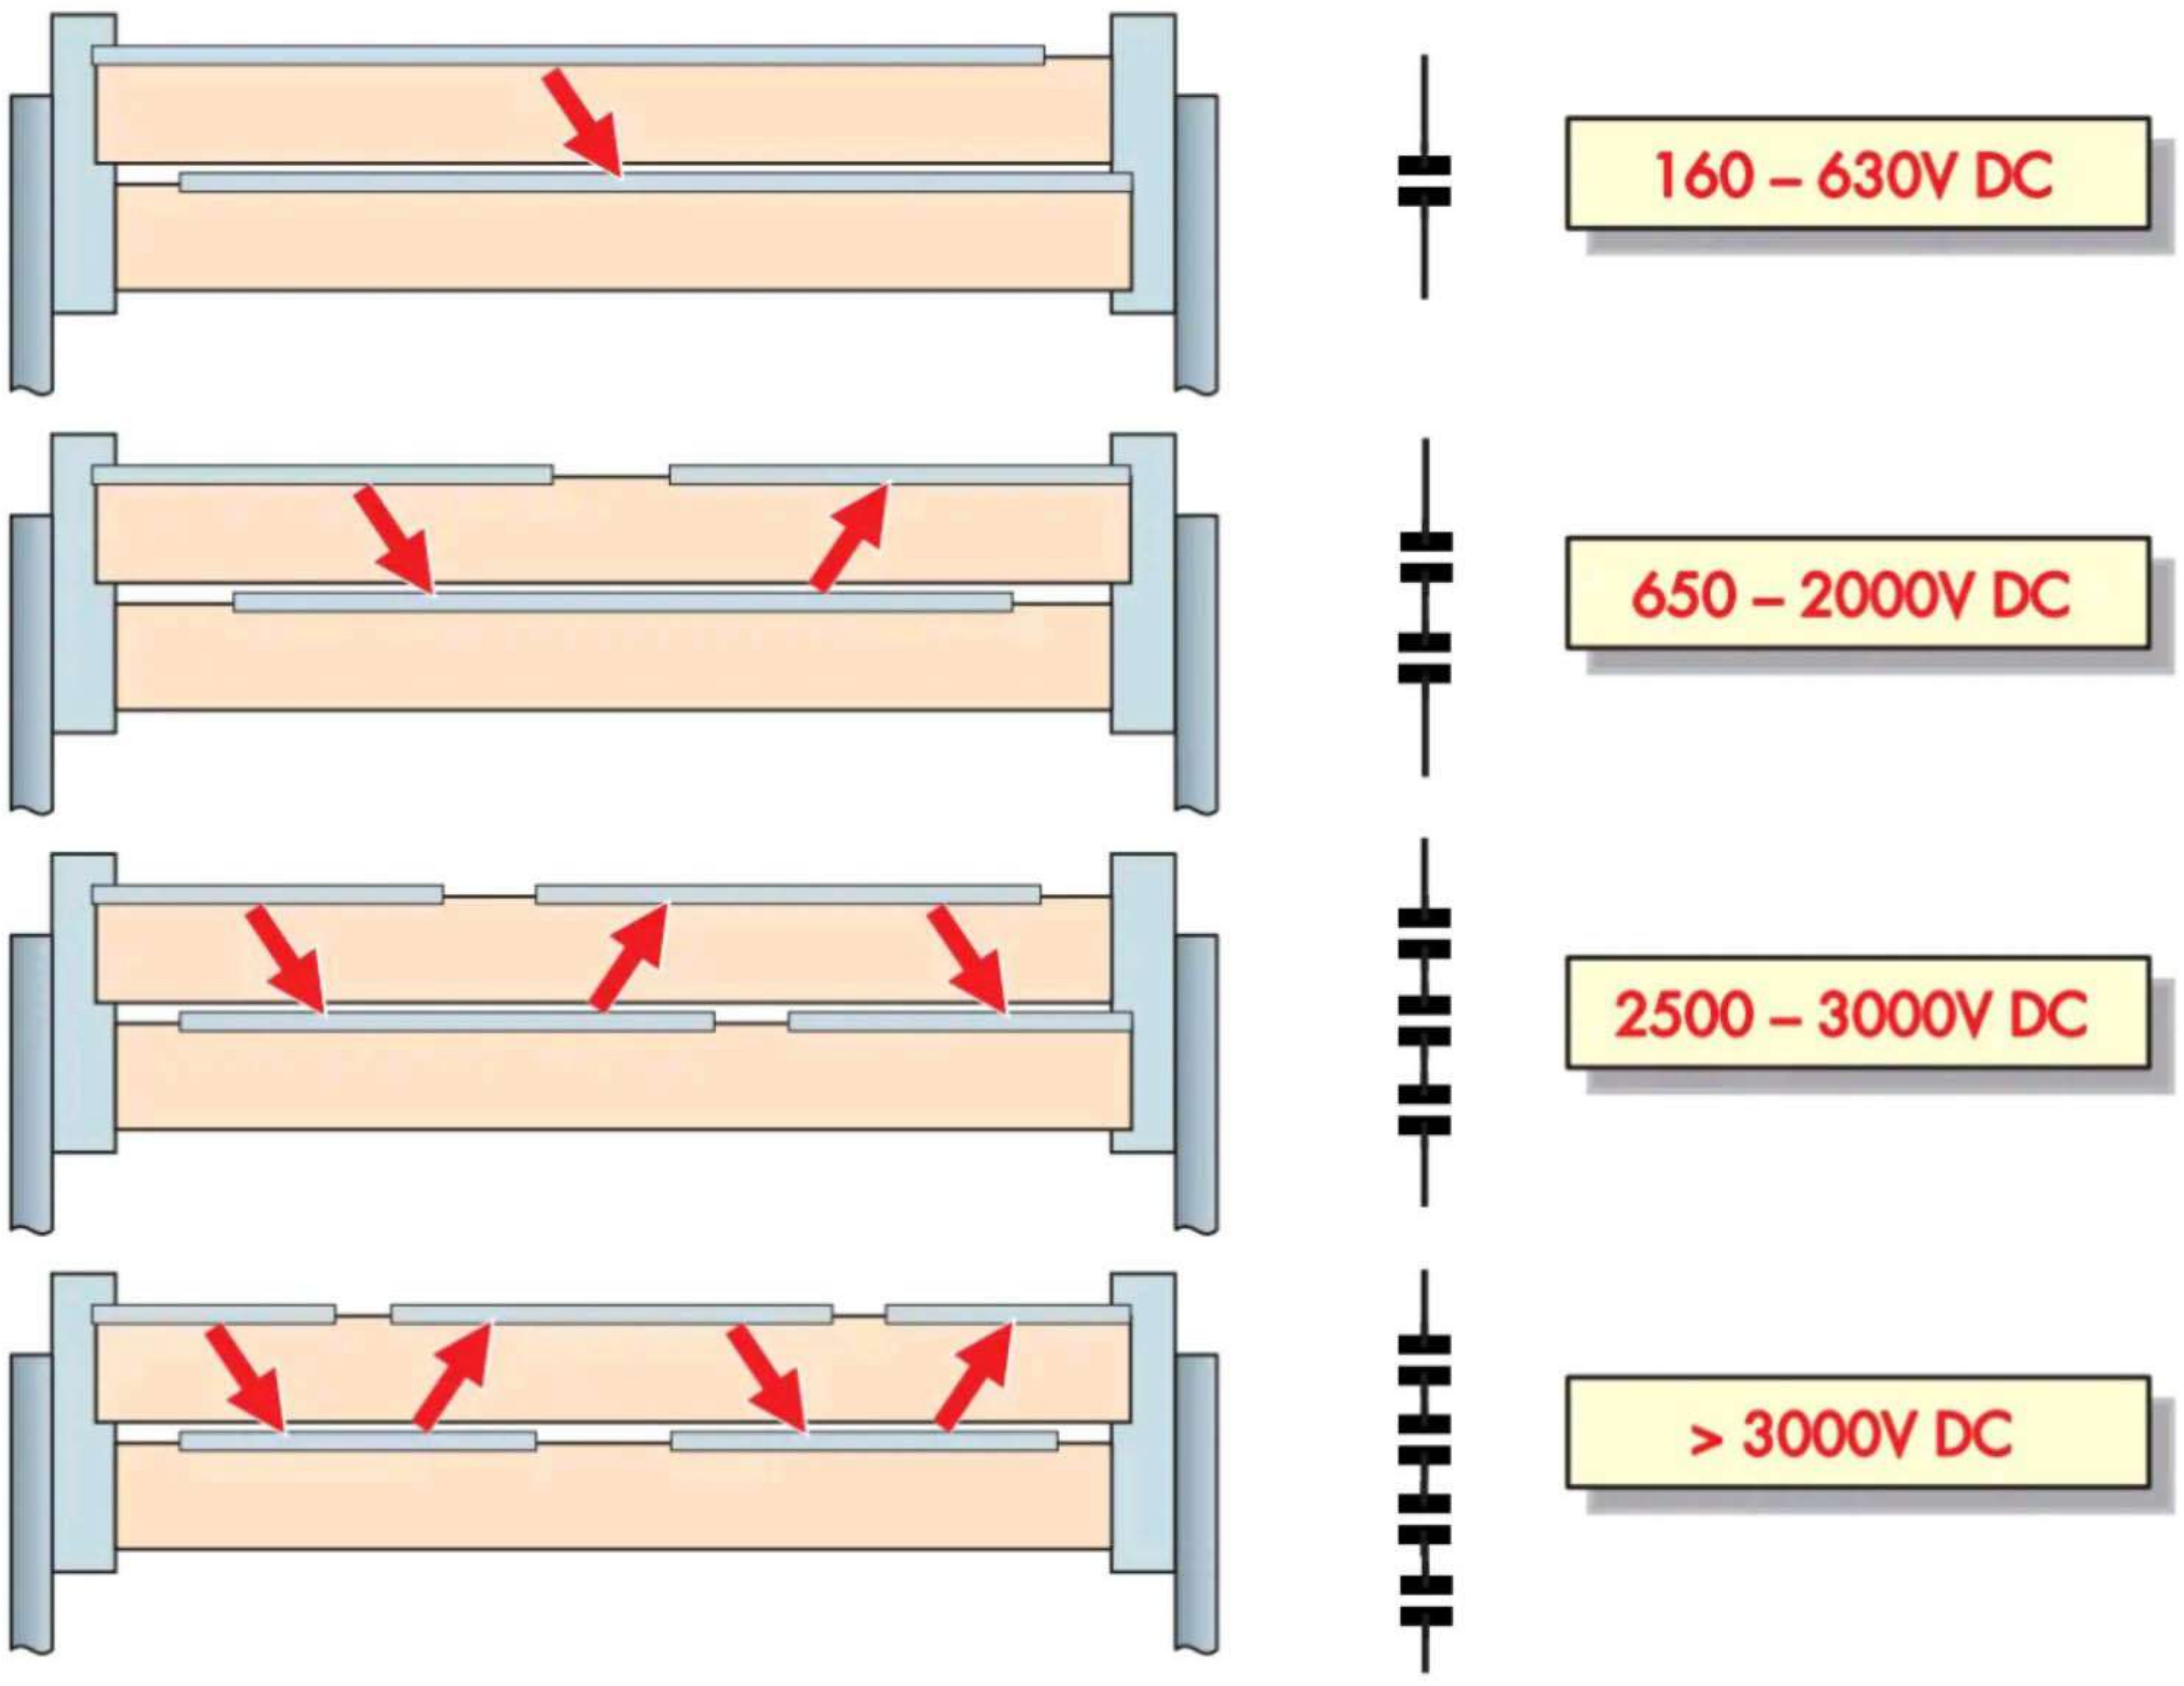

Fig.20: the basic method of plastic film capacitor manufacturing doesn't work very well for applied voltages above about 630V. Higher voltage ratings are possible, but the capacitor needs to be internally separated into several elements connected in series, so that the dielectric material only has a fraction of the applied voltage across it. *(original source: Wikipedia)*

A metallisation layer is sprayed onto the ends of the roll to connect the layers, and leads are attached. This is why the conductors are offset; each layer is only exposed at one end of the roll, or else the sprayed metal layer will short out the capacitor. After this, the capacitor is typically impregnated with silicone oil or another insulating fluid to prevent moisture ingress.

The terminals are then attached and the capacitor is encapsulated, sometimes by being dipped, other times by being sealed into a pre-formed plastic case.

To make plastic film capacitors with a voltage rating above 630V DC, partial metallisation can be used to effectively form multiple capacitors in series using the same basic techniques. This can extend voltage ratings up past 3000V DC (see Fig.20).

As well as small-signal capacitors and those for mains filtering, plastic dielectric capacitors are also used in motor run applications. Many motor start capacitors are electrolytic types, but electros are not suitable for handling the continuous current and high voltages that motor run capacitors are subjected to.

So they are typically made with polypropylene or similar plastic dielectrics and thick metal films to handle high currents continuously.

Electrical double-layer capacitors (EDLs)

These are often known as supercapacitors or ultracapacitors. They are a variation on electrolytic capacitors, with extremely high capacitances but usually low voltage ratings, and often very high internal resistances (and thus low operating currents). Silicon Chip published an article on ultracapacitors in April 2008 (siliconchip.au/Article/1793).

EDLs use similar conductive polymers with a very high surface area for both the positive and negative electrodes, with a common electrolyte in contact with both. Anions and cations in the electrolyte form insulating Helmholtz layers in direct contact with the surfaces of both electrodes.

These layers are only one atom thick, and as mentioned at the start of the article, capacitance is proportional to

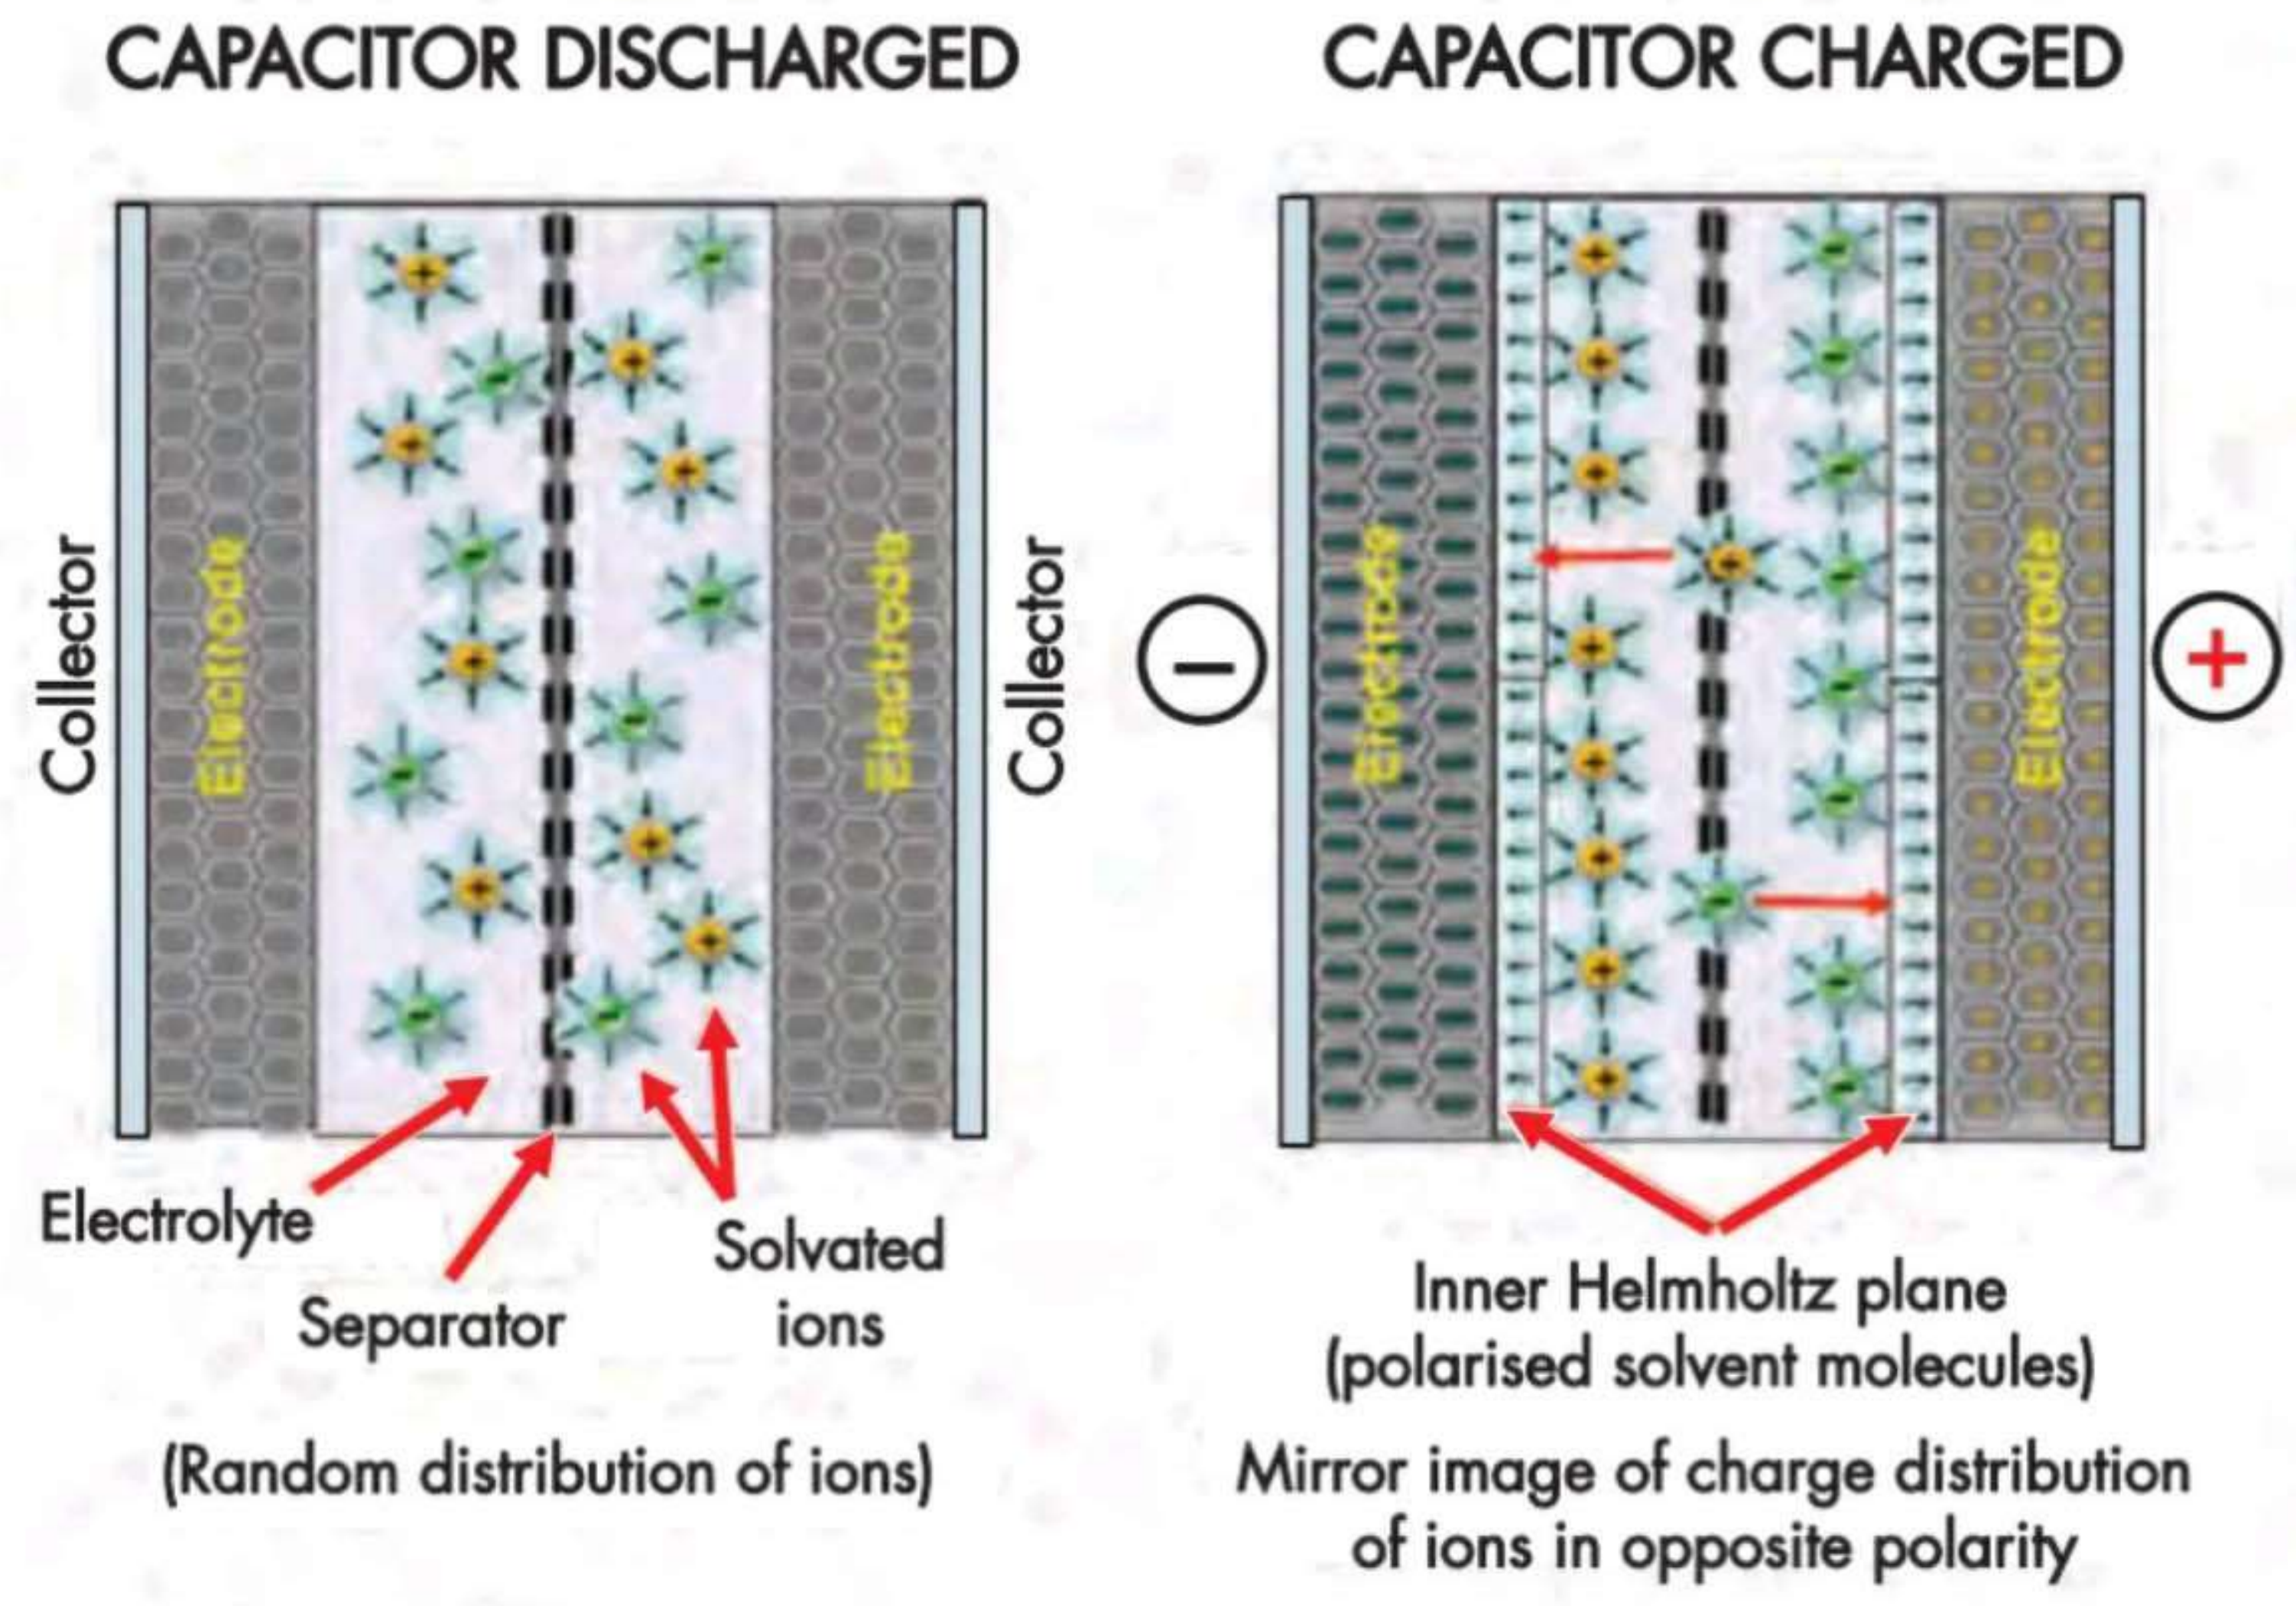
CAPACITOR DISCHARGED
CAPACITOR CHARGED
Collector
Electrode
Electrode
Collector
−
Electrode
Electrode
+
Electrolyte
Separator
Solvated ions
(Random distribution of ions)
Inner Helmholtz plane (polarised solvent molecules)
Mirror image of charge distribution of ions in opposite polarity

Fig.21: as the name suggests, a double-layer (EDL) capacitor effectively has two dielectric layers, one at the surface of the anode and one at the cathode, with a conductive electrolyte between the two. The advantage is that these layers are super-thin, just one molecule wide, giving extremely high capacitances in a small package. However, this thin dielectric layer results in a very low voltage rating, typically either 2.7V or 5.5V (source: Wikimedia user Elcap).

712
NEC5.5
FS 0.047F
NEC/TOKIN
FYD 5.5V
1.0F

surface area and inversely proportional to dielectric layer thickness. You can't get a much thinner layer than one atom (see Fig.21).

Given the large surface area of the electrodes, EDL capacitors can have values exceeding one Farad in a package not much bigger than a can 19mm in diameter and about 16mm tall.

The fact that the current must pass through two polymer layers plus an electrolyte, neither of which is especially conductive, is why the current delivery of EDLs is generally limited. The extremely thin dielectric layer is the reason why voltage ratings of only 2.7V or 5.5V are common.

Both of these problems can be mitigated by connecting many EDL capacitors in parallel (to improve current handling) or series (to increase the voltage rating, at the expense of capacitance). Higher voltage EDLs usually have multiple internal EDLs in series.

You might be using an ultracapacitor without realising it; Mazda introduced its i-ELOOP system in vehicles from 2011, and it is now in many vehicles. This system recovers kinetic energy during braking to rapidly charge an ultracapacitor, then uses that energy to charge the vehicle battery over a longer period.

Other types of capacitor

Silvered mica capacitors unsurprisingly use mica, a type of mineral, as the dielectric. Mica was chosen both for its good dielectric properties and because its crystalline structure makes it very easy to cleave into super-thin sheets; just what you need to achieve a decent capacitance. A thin layer of silver is applied to each side, and voila, you have a capacitor with excellent linearity and low leakage.

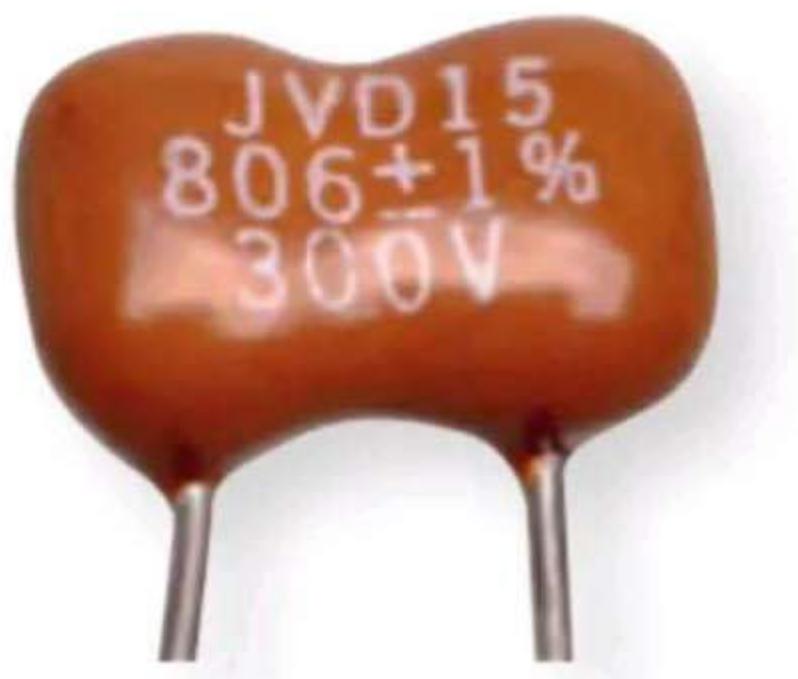
JVD15
806±1%
300V

An example of an 806pF 300V mica capacitor. The 1% rating means its actual value will be in the range of ~798-814pF.

Mica capacitors have mostly been replaced by ceramic or plastic film types, as both are significantly cheaper to manufacture and achieve similar performance.

Some still value mica caps for audio circuits. Besides good linearity, another property of mica capacitors is that they usually have tight tolerances due to their predictable thickness, measurable surface area and low temperature coefficient.

Another type of non-polarised capacitor that was widely used but is now far less common (although still available) is the paper capacitor (sometimes known as an MP [metallised paper] capacitor).

These have also mostly been supplanted by ceramic or plastic film capacitors. The main disadvantage of paper capacitors is that they can absorb moisture from the air and fail; older types have also been known to catch fire! Modern capacitors usually combine paper and plastic (usually PET or polypropylene) to overcome these disadvantages. Their main advantage is low cost.

One benefit that paper capacitors retain is that they usually have zinc metallisation compared to the aluminium metallisation of plastic capacitors. This provides better 'self-healing' capabilities due to its lower-energy evaporation process.

Variable capacitors

Variable capacitors work either by varying the amount of overlap between two sets of metal plates, or by chang-

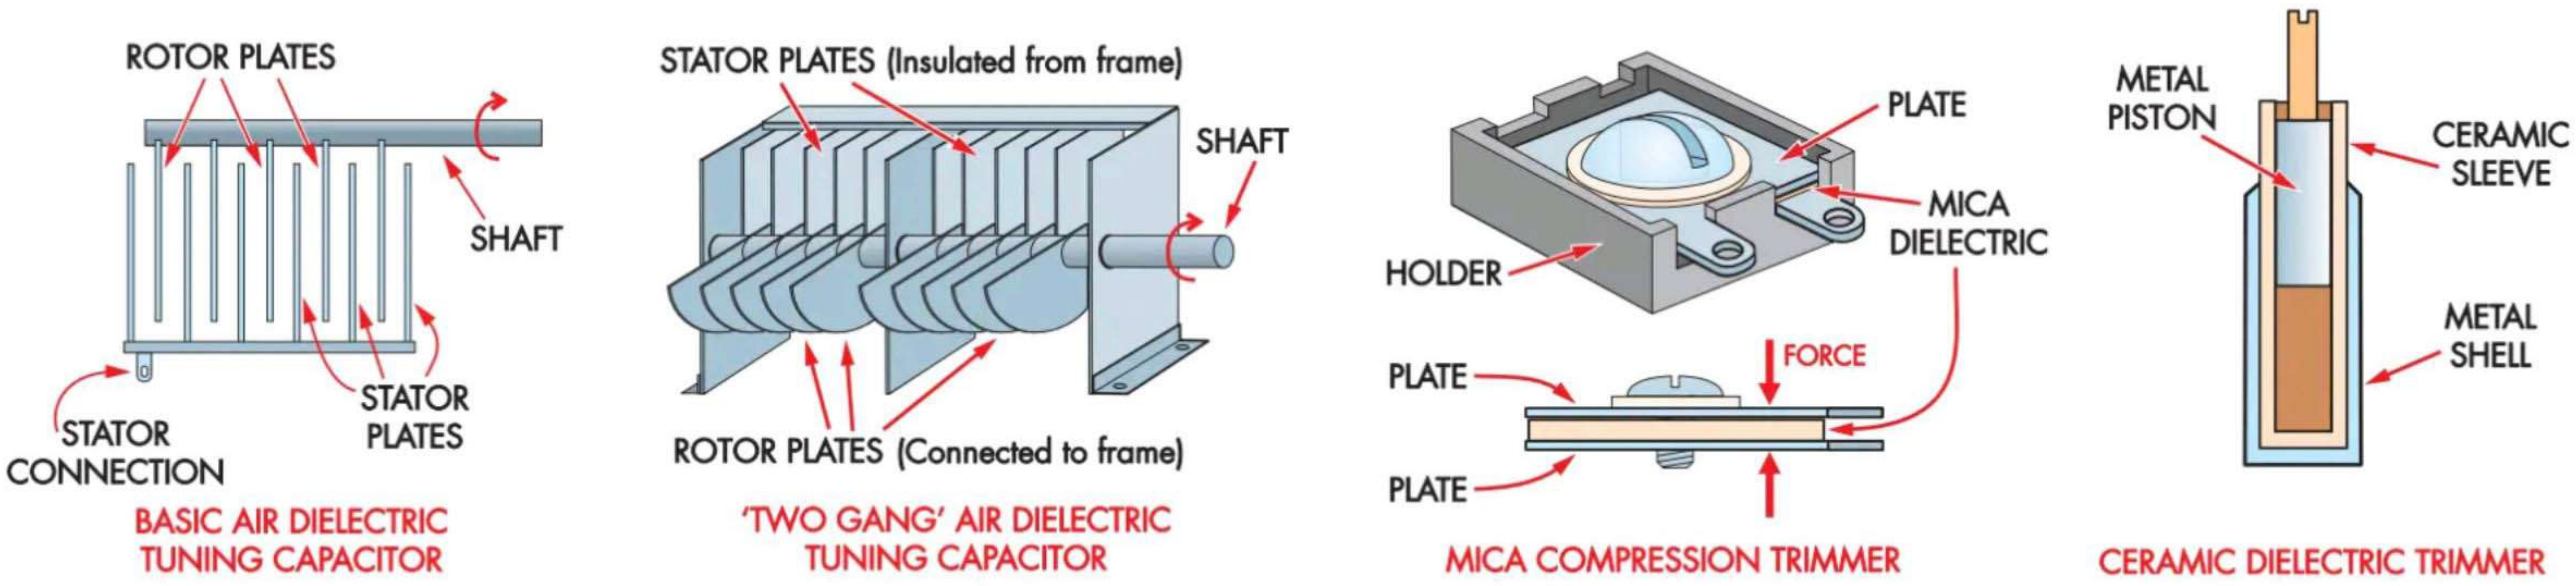
ROTOR PLATES
SHAFT
STATOR PLATES
STATOR CONNECTION
BASIC AIR DIELECTRIC TUNING CAPACITOR
STATOR PLATES (Insulated from frame)
SHAFT
ROTOR PLATES (Connected to frame)
'TWO GANG' AIR DIELECTRIC TUNING CAPACITOR
PLATE
MICA DIELECTRIC
HOLDER
PLATE
FORCE
PLATE
MICA COMPRESSION TRIMMER
METAL PISTON
CERAMIC SLEEVE
METAL SHELL
CERAMIC DIELECTRIC TRIMMER

Fig.22: the simplest type of variable capacitor, used in many vintage radios, is just two sets of interleaved metal plates where the amount of overlap can be adjusted. Air is the dielectric. Miniature trimmer caps tend to use a plastic or mica dielectric and bring the two plates closer together or further away to vary the capacitance.

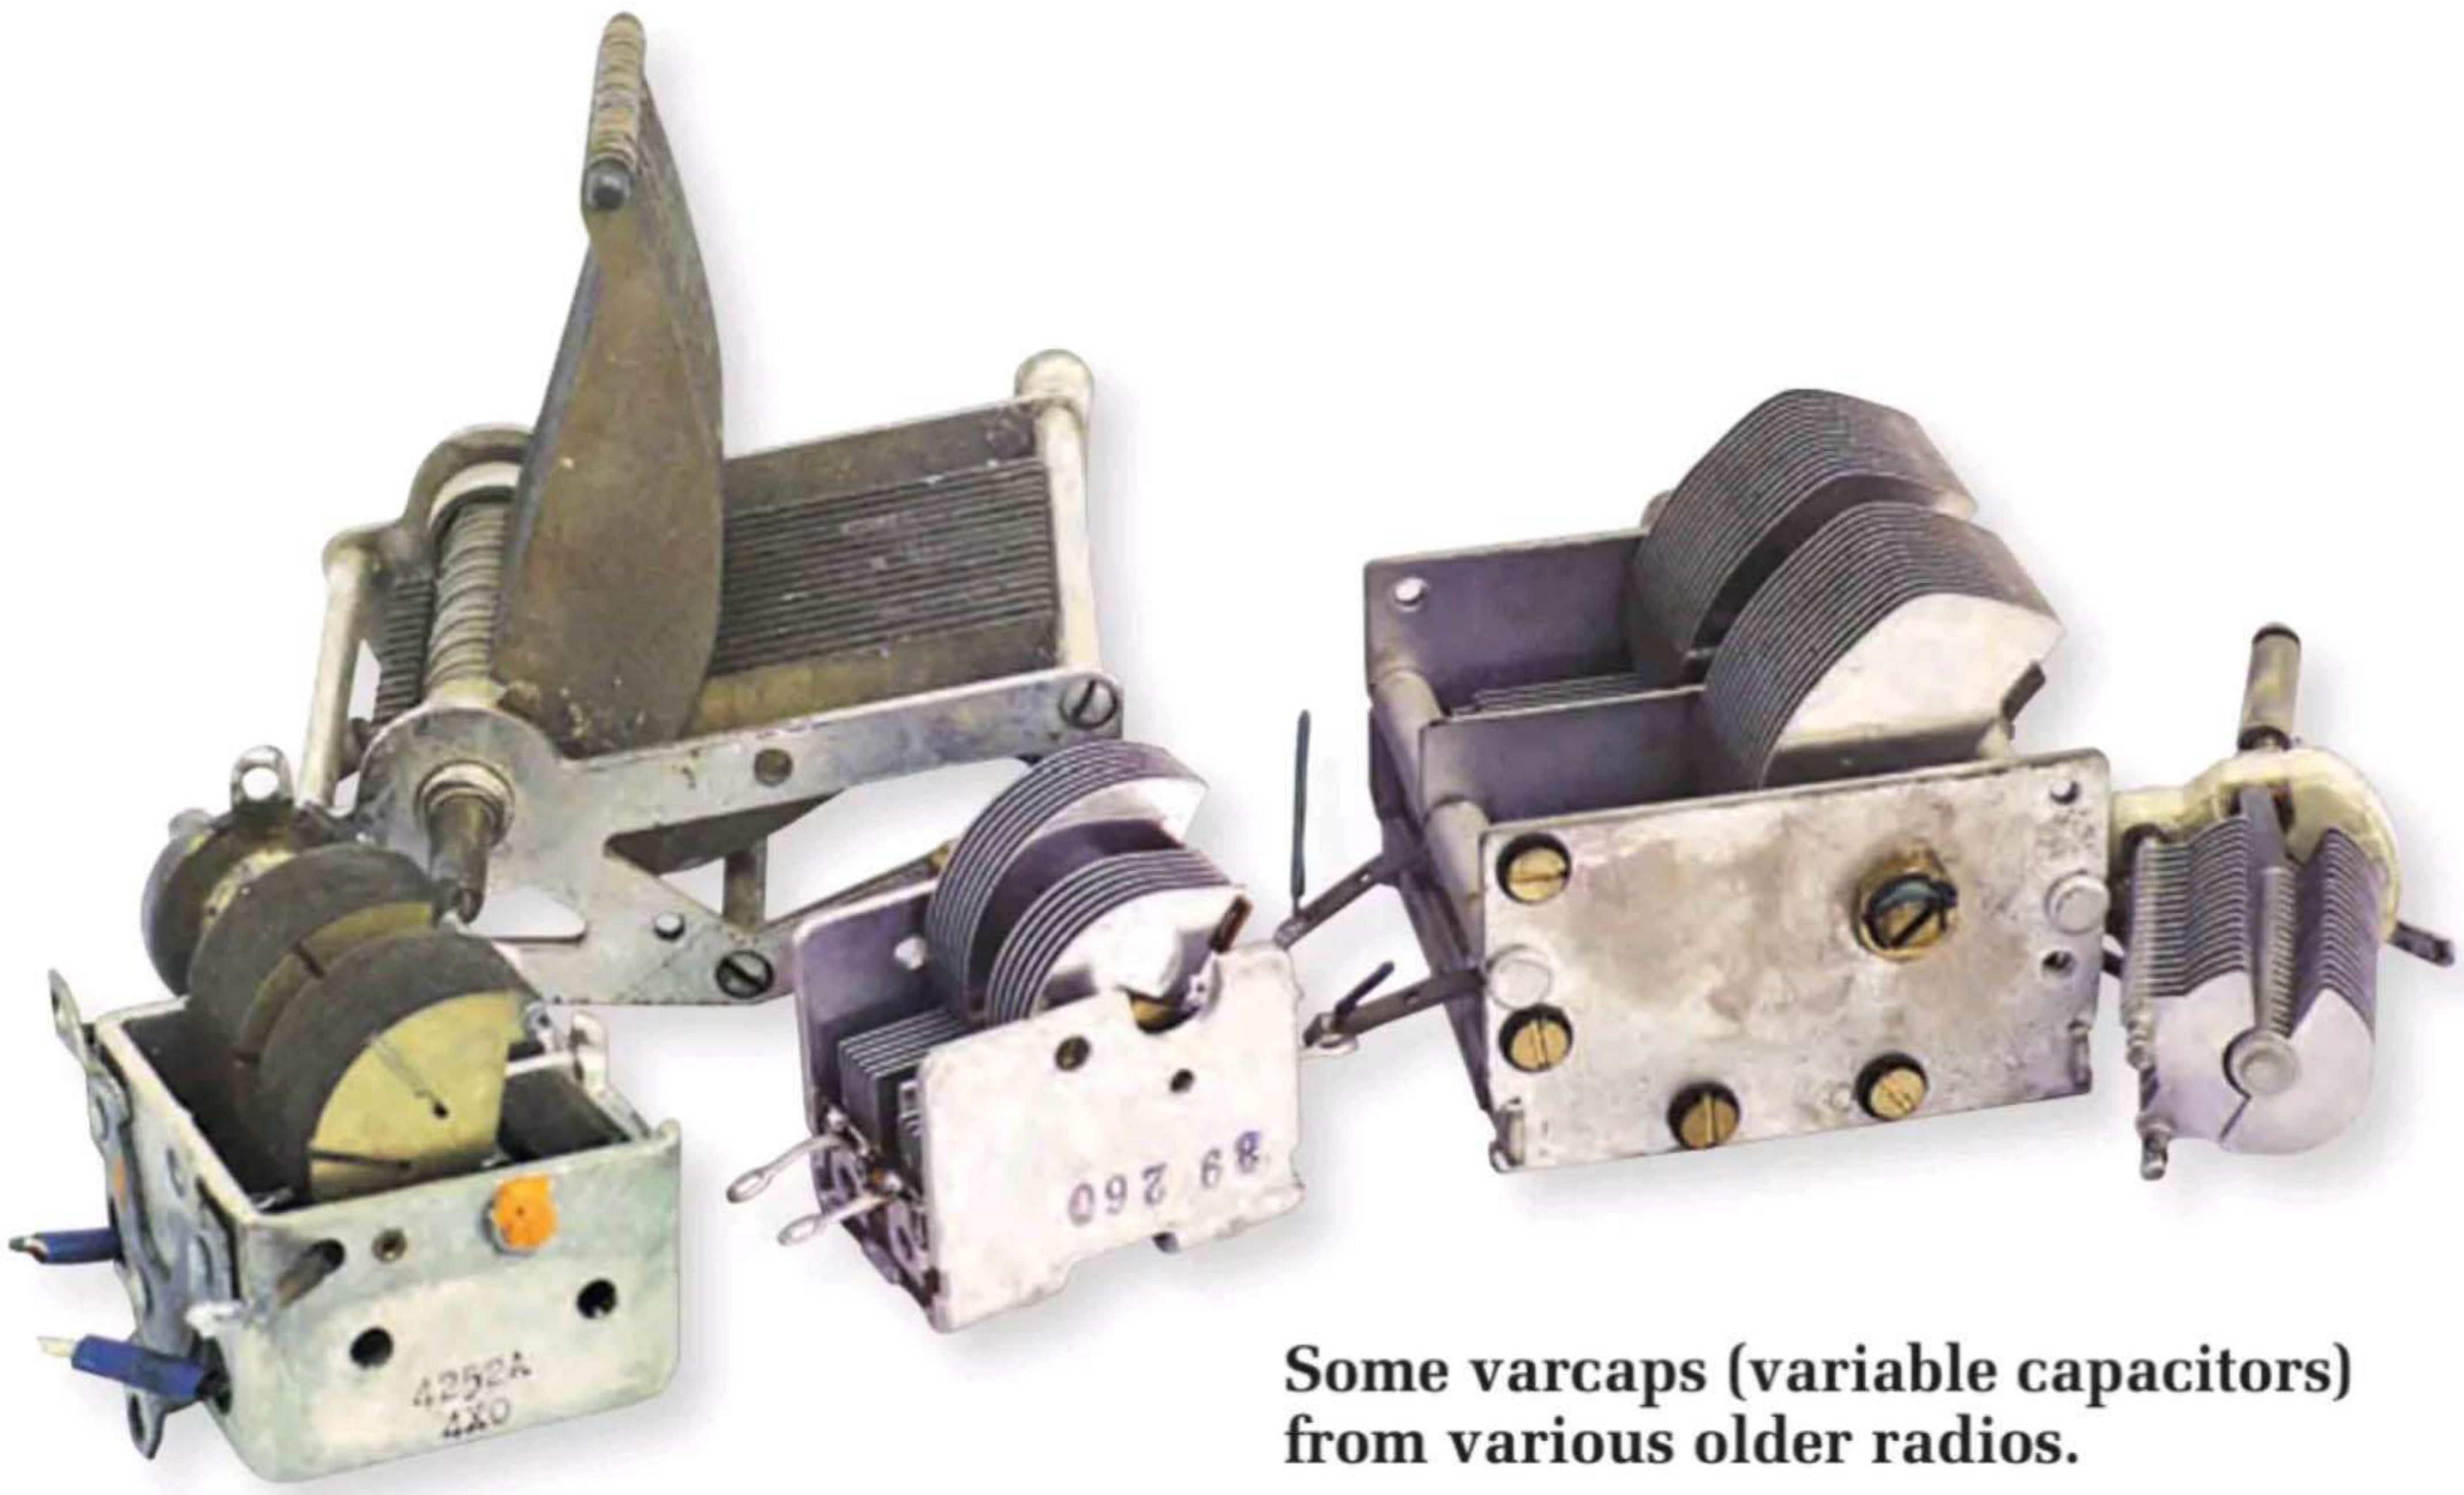

Some varcaps (variable capacitors) from various older radios.

ing the spacing between two plates (possibly separated by a plastic or mica dielectric). In many cases, the dielectric is air.

As such, the range of capacitance for a typical tuning capacitor is generally from a few picofarads up to a few hundred picofarads (see Fig.22).

They are typically used as part of an RF oscillator or filter circuit, so low picofarad values give the required time constants with reasonable values for other components (usually inductors).

Oddball types

The capacitors described above probably cover 99% of the capacitors you might come across, but other types exist. That includes those with a glass or silicon dielectric, or even a vacuum!

Capacitor parameters

In addition to its construction/dielectric (ceramic, aluminium electrolytic, tantalum electrolytic, plastic film, plastic foil etc), a capacitor is described by its capacitance, tolerance, voltage rating(s), ripple current rating(s), leakage current rating(s), operating temperature and expected lifespan.

Furthermore, each capacitor type has several associated performance metrics, which may be fixed or vary with parameters like temperature, applied voltage, signal frequency, age etc. These include the ESR (equivalent series resistance), ESL (equivalent series inductance), dissipation factor (DF or delta [Δ]), temperature coefficient (tempco), voltage coefficient, linearity and more.

We'll describe all of these, starting with the parameters which typically form part of a model or part code.

Capacitance: the nominal capacity of the device, measured with no or little charge (typically around 0.5-1.0V across the capacitor) and at room temperature.

For very low-value capacitors (fraction of a picofarad to a few picofarads), the measured capacitance can be affected by the connected PCB tracks/pad or wires, or even the device's lead length.

Tolerance: how close you can expect the capacitance to be compared to the nominal value. If you have a 10μF ±10% capacitor, if its value is less than 9μF or more than 11μF, then it would be considered faulty.

However, during actual use, its capacitance could vary outside this range, as explained below.

Voltage rating(s): the applied DC voltage across the capacitor terminals can safely vary from 0V up to this figure. For non-polarised types like ceramic or plastic film, it can also be negative, meaning the full range of operating voltages is effectively doubled (ie, -50V to +50V for a 50V capacitor).

Some capacitors have a higher 'surge' voltage rating which will not damage them if applied for a limited period, but that is less common these days. Note that you sometimes need to keep the voltage below the rating for good performance; more on that shortly.

Ripple current rating(s): all capacitors have some intrinsic resistance and therefore heat up as current passes through them; current flows through a capacitor during both charging and discharging.

For example, if a capacitor is used to filter the output from a bridge rectifier turning AC to DC, it supplies the full load current most of the time (when the bridge is not conducting). But it also must absorb large pulses of current to recharge when the bridge comes into conduction, 50 or 100 times per second.

Such circuits must be designed to avoid exceeding the RMS ripple current rating of the filter capacitor(s) or else they can rapidly overheat and fail.

There are generally different figures given at low (50/100Hz) and high (100kHz) frequencies, due to the changing impedance of the capacitor, from both its capacitance and its ESL (see below). The ripple current rating is usually higher at higher frequencies.

Leakage current rating(s): some capacitors (eg, electros) can have a fairly significant leakage current through the capacitor even when the voltage is steady. This is usually proportional to the applied voltage.

Ceramic, plastic film and mica capacitors also have leakage currents, although they are generally very low and are often (but not always) of no concern. This is important in some applications, like sample-and-hold buffers, where the voltage across a capacitor must remain stable for relatively long periods.

Operating temperature: critically for electrolytic capacitors, this is the maximum temperature at which the capacitor is guaranteed to meet the stated performance figures. It is also the temperature at which the expected lifespan (if given) is calculated.

Capacitor lifespan roughly doubles for each 10°C below the rated temperature, and halves for each 10°C above it. Typical ratings are 85°C, 105°C and 125°C. We recommend using 105°C-rated capacitors with an expected lifespan of at least a few thousand hours to avoid early failures.

Lifespan: usually stated in thousands of hours MTBF (mean time between failures), with 1000 hours at the lower end and about 10,000 hours at the upper end. If you can find a capacitor rated to last for 10,000 hours at 125°C, it'll probably outlast the rest of the circuit!

Performance metrics

Equivalent series resistance (ESR): this is a crucial metric for most capacitors as it has a strong effect on how well the capacitor can smooth DC voltages, and how much heat is generated at higher currents.

You can think of a real capacitor like an ideal capacitor with a resistor in series; the lower the value of that resistor, the less it is 'isolated' from the circuit it is connected to.

Typical ESR values are a few ohms for a low-value electrolytic capacitor, down to a few milliohms for a large, low-ESR electrolytic, tantalum, polymer or ceramic capacitor.

Equivalent series inductance (ESL): just like ESR, you can imagine that all capacitors have a low-value inductor internally connected in series with the capacitance.

This has little effect at low frequencies, but can make the effective impedance of the capacitor so high that it is useless at higher frequencies.

ESL is critical for applications like bypassing multi-GHz ICs such as CPUs and RF devices. Smaller capacitors generally have a lower ESL, and certain construction methods can dramatically reduce the ESL of larger capacitors.

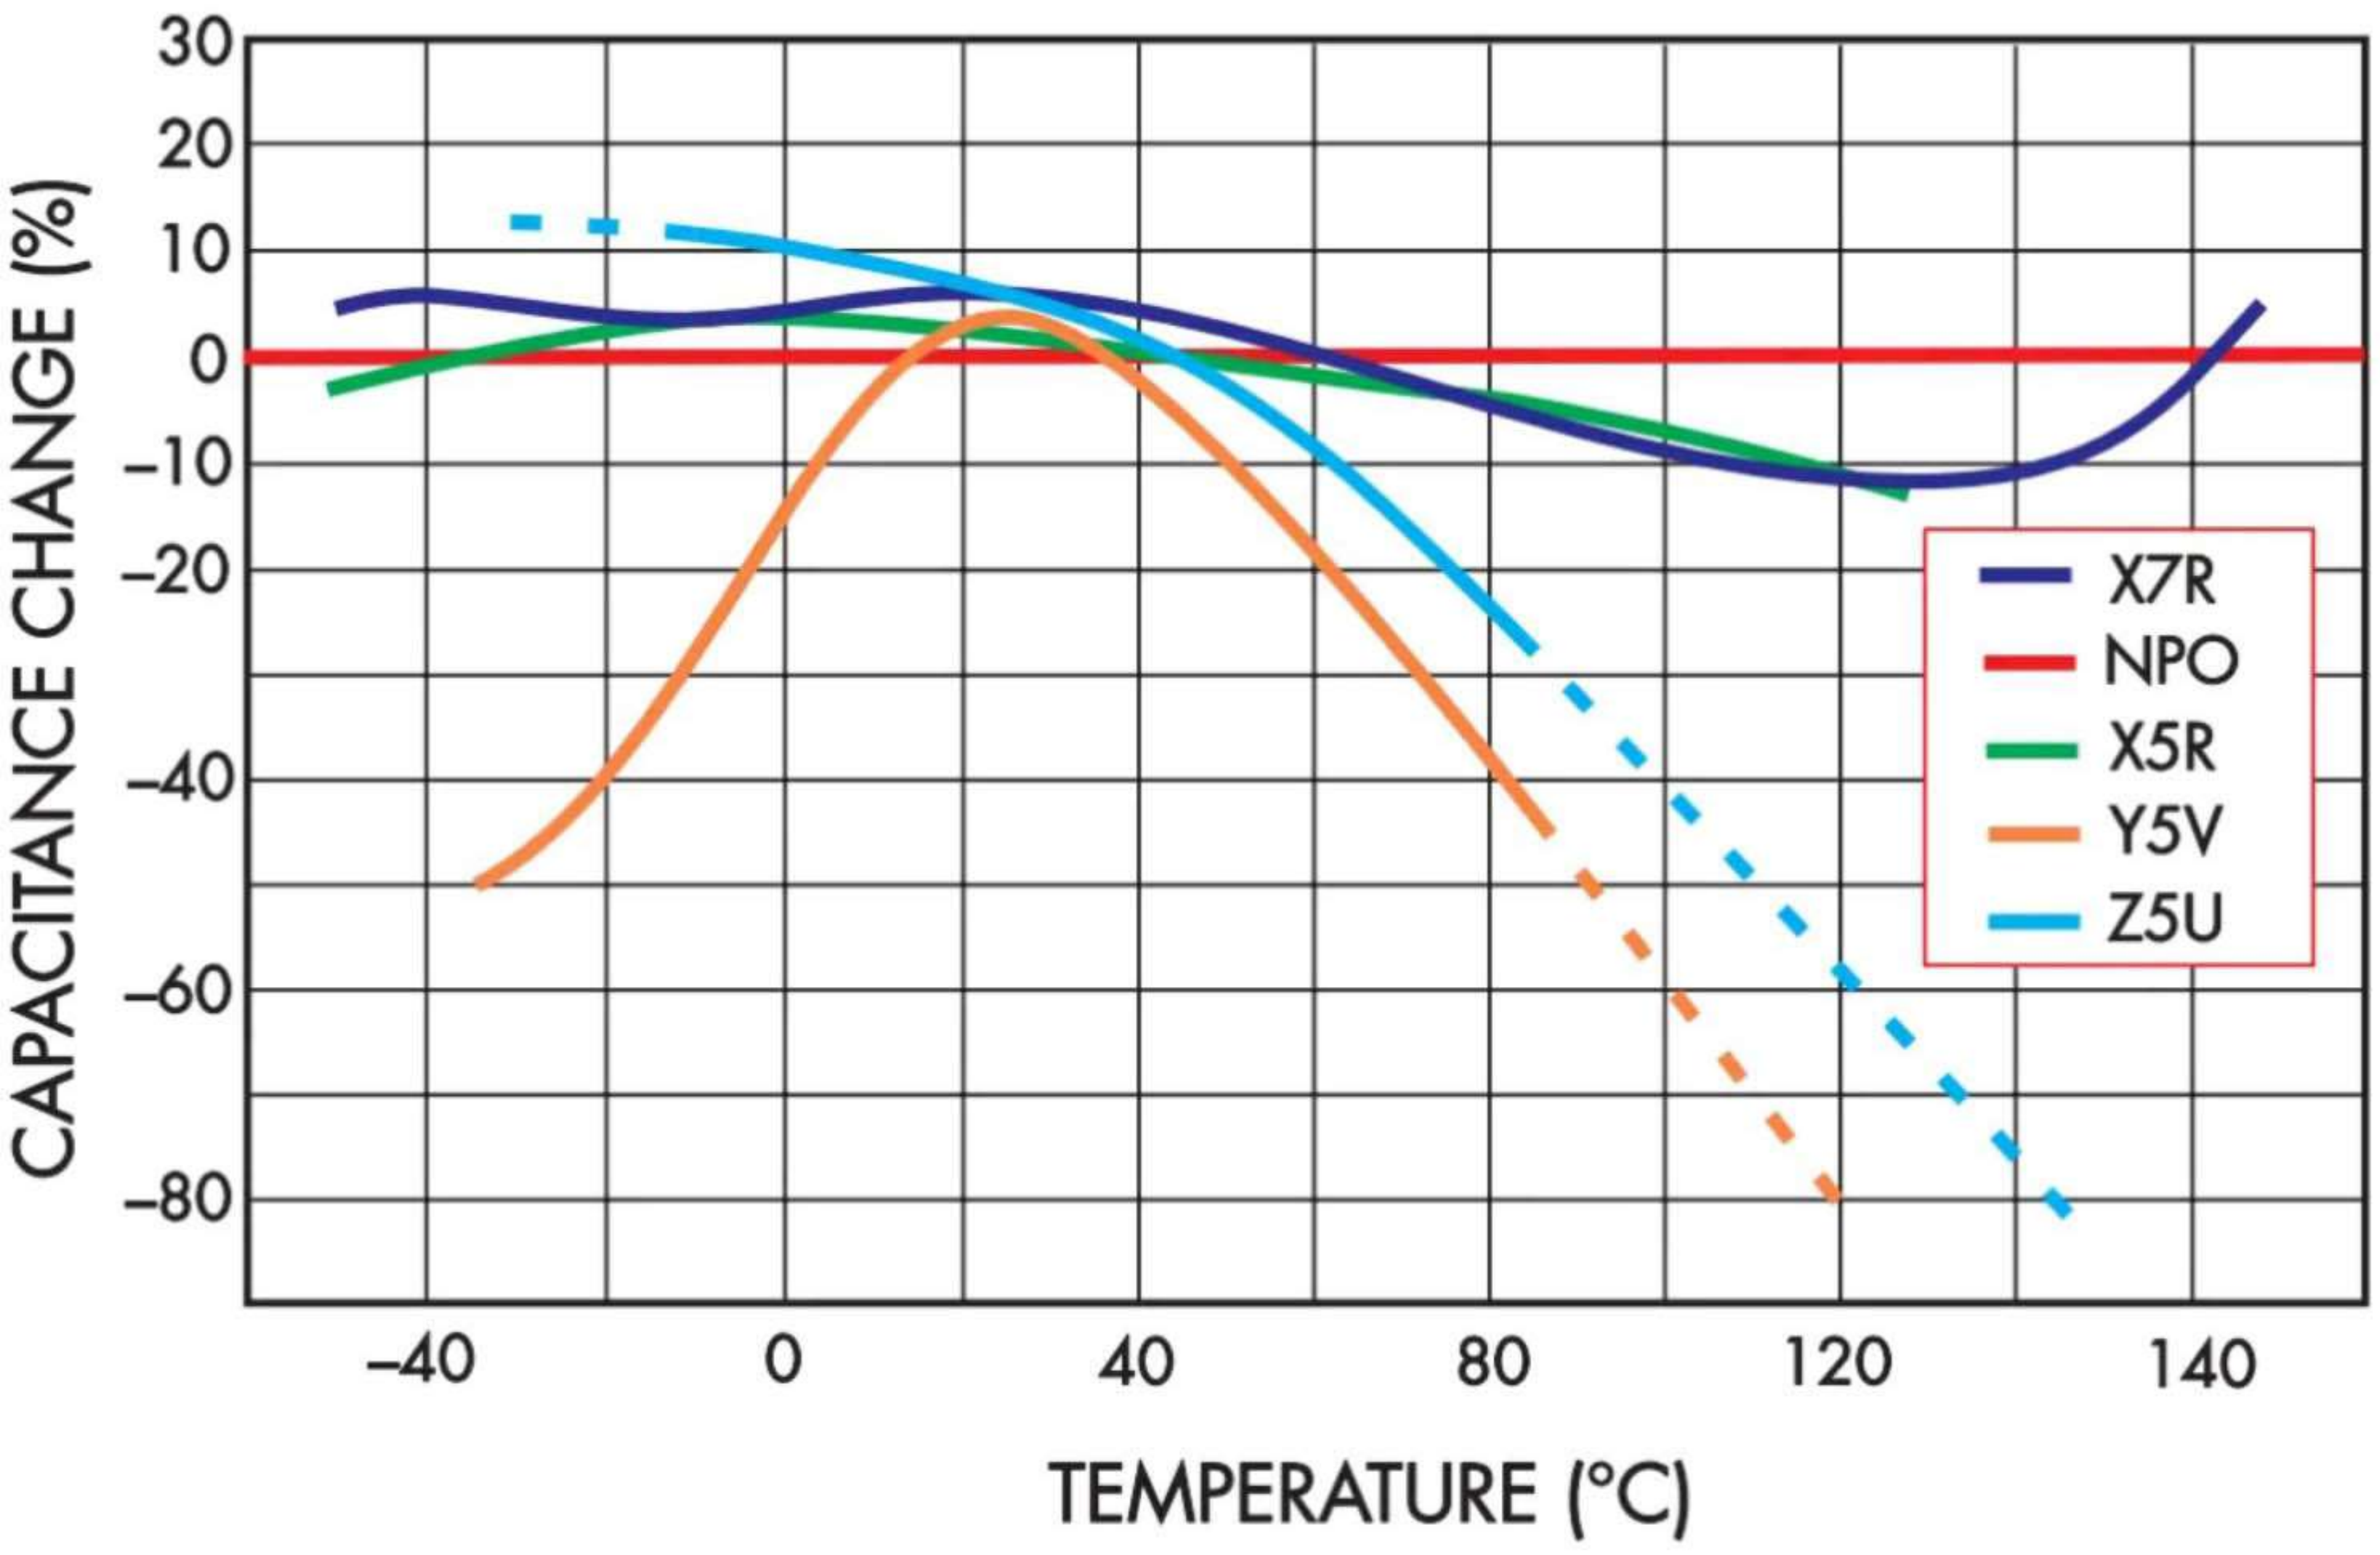

Fig.23: by definition, the temperature coefficient of an NP0 ceramic capacitor is zero (or very close to it). On the other hand, Y5V ceramic capacitors vary in value wildly with temperature. Z5U is a little better at lower temperatures, but still poor at high temperatures. X5R and X7R are the 'go to' ceramic dielectrics because they are cheaper and more compact than NP0 capacitors, but have a much more modest temperate coefficient than Y5V or Z5U. *(source: Wikipedia & Johanson Dielectrics)*

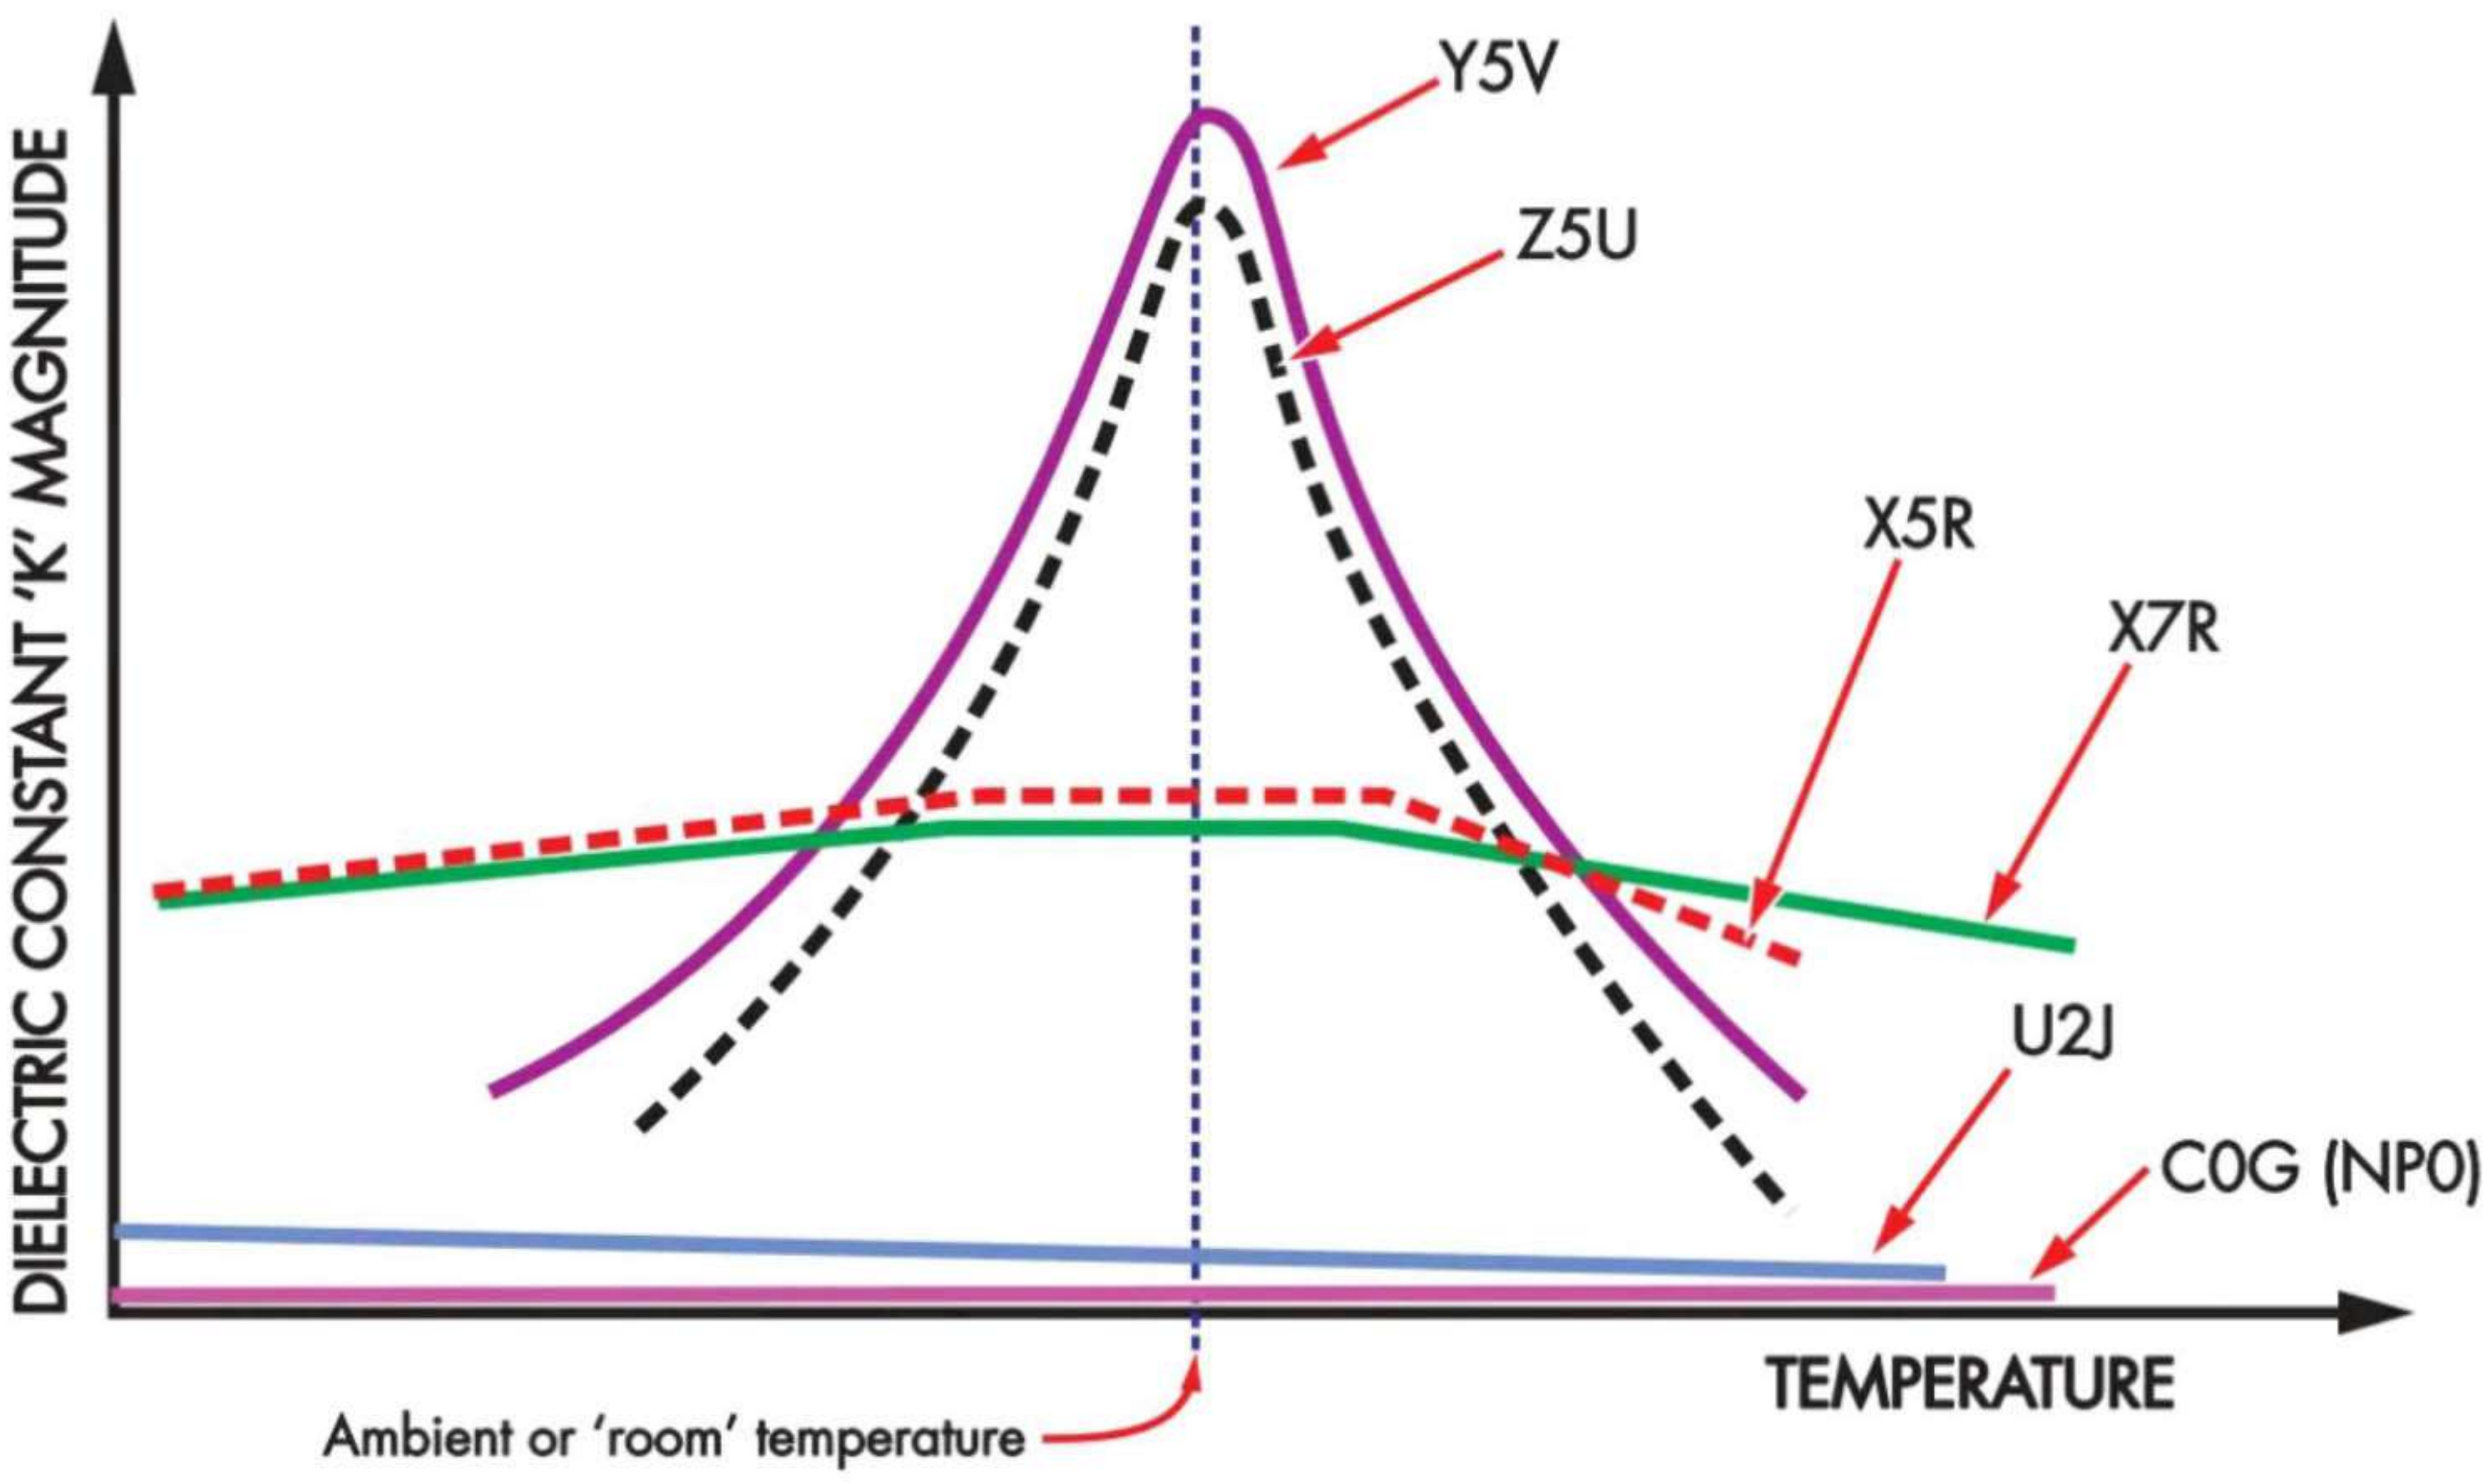

Fig.24: the variation with temperature of the dielectric constant, K, for several ceramic materials. You can see that X5R and X7R have a much higher K than C0G/NP0, making for higher capacitances in a smaller volume, with only a slight variation over the temperature range. The vast variation for Y5V and Z5U makes them unattractive. While they give a high capacitance at room temperature, at very high or low temperatures, the K value drops below that of both X5R and X7R. *(original source: Digi-Key)*

The combination of capacitive reactance, ESR and ESL produces a characteristic valley-like impedance curve for most capacitors, as shown previously in Fig.10.

Dissipation factor (DF): also known as $\tan(\delta)$, is the reciprocal of the ratio of the ESR and capacitive reactance, and as such, is typically a number close to but slightly less than one.

It can be easier to tell how close to ideal a given capacitor is by looking at the DF rather than ESR, and it can also make it easier to compare the performance of capacitors with different values of capacitance.

Temperature coefficient (tempco): how much the capacitance changes with temperature. In an application where the capacitance is used to define a frequency (eg, in an oscillator or filter), this must be minimised to prevent frequency drift with temperature. Hence, NP0/C0G ceramics or plastic film capacitors are generally preferred in those roles (see Figs.23 & 24).

This is very important to keep in mind with high-K dielectrics like Y5V ceramics; they can lose 80% or more of their capacitance at elevated temperatures! That's made even worse by…

Voltage coefficient: how much the capacitance changes as the capacitor is charged.

This causes two main problems. One is a loss of effective capacitance; combined with the poor tempco of Y5V ceramics, a 10µF 6.3V Y5V capacitor at 80°C, charged to 5V, might have less effective capacitance than a 1µF 50V X5R capacitor under the same conditions! (See Figs.25 & 26). This is why we steer well clear of cheaper Y5V ceramics.

The other problem with a high voltage coefficient is that it dramatically impacts linearity.

So, perhaps unintuitively, standard aluminium electrolytic capacitors are better for audio coupling than most high-quality multi-layer ceramic capacitors. NP0/C0G ceramics are the exception, but they are huge and horrendously expensive at the sort of values required for most audio coupling applications.

Linearity: this is not something you will find on most data sheets, and there is no standard way of representing it. But it is a real effect that varies significantly between different capacitor types.

The easiest way to measure it is to form a simple RC filter (low-pass or high-pass) with a relatively low-value, linear (thin film) resistor and a capacitor. You then feed a very pure sinewave into the filter, at a frequency near the -3dB point, and measure the distortion figure of the voltage across the capacitor.

The resulting % THD is inversely proportional to the capacitor's linearity. A very linear capacitor like a polypropylene or NP0/C0G ceramic capacitor will introduce an unmeasurable level of distortion (below 0.0001%). Other plastic film types like polyester are slightly worse, resulting in a measurable but not worrying level of distortion (say 0.0005%).

Other capacitors like high-K ceramics, electrolytics and so on could give distortion measurements of 1% or more, reflecting the fact that their I/V curves are not straight.

In some cases (eg, electros), the curves can even have hysteresis, meaning they are a different shape for charging and discharging.

This is most important in audio circuits, although other circuits (eg, RF) might be sensitive to linearity too. Note that electros are fine for audio coupling, even though they are not terribly linear, as the applied AC voltage in that role is so small that it isn't a significant effect (unlike the voltage coefficient of many ceramics, which makes them unsuitable for that role, despite probably being more linear than electros).

Ageing: capacitors can change value over time, usually decreasing due to degradation of the dielectric. Typically, those with a tighter initial tolerance will tend to maintain their capacitance better over time (see Fig.27).

This is apart from an actual failure of the component, which might manifest as a much lower capacitance, higher ESR, higher leakage (especially at voltages approaching the rating) or some combination of the three.

Further reading

- **Types:** https://w.wiki/q86
- **Ceramic:** https://w.wiki/q87
- **Electrolytic:** https://w.wiki/q88
- **Tantalum:** https://w.wiki/q89
- **Polymer:** https://w.wiki/q8A
- **Film:** https://w.wiki/q8B
- **Supercaps:** https://w.wiki/q8C PE

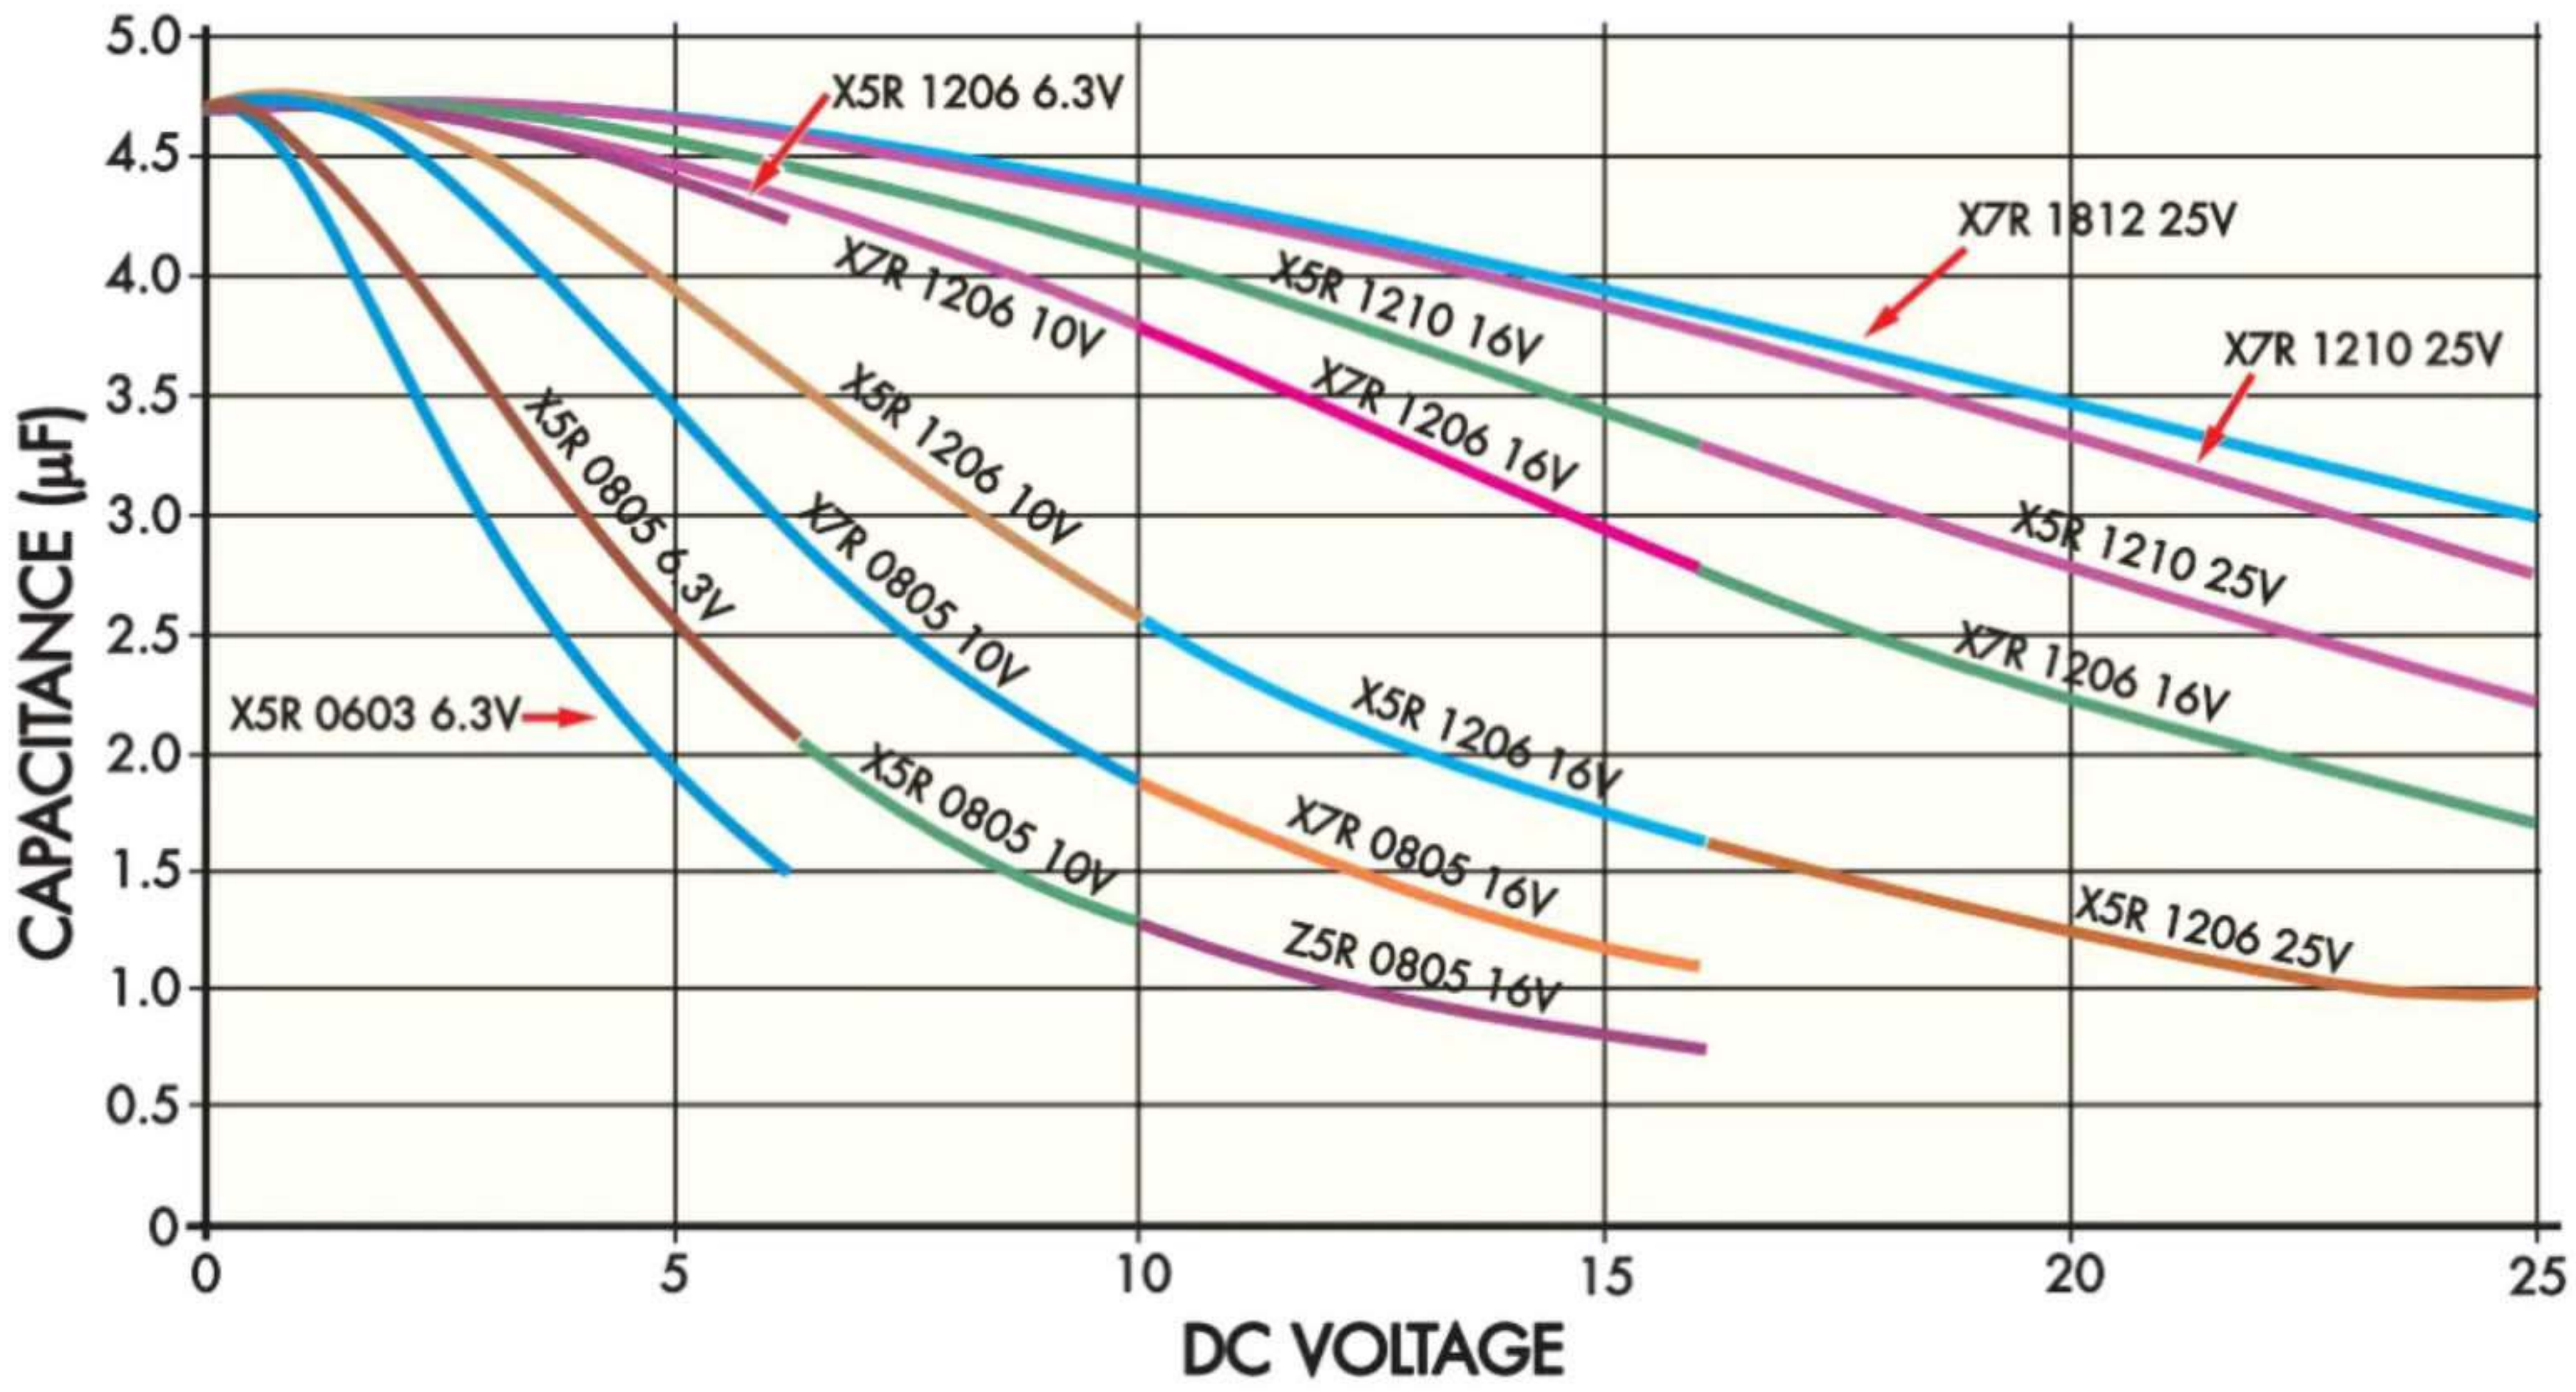

Fig.25: the measured value of a range of 4.7µF X5R and X7R capacitors with the application of a range of DC voltages. Note how the physically larger capacitors tend to retain their capacitance better as they are charged to a similar voltage. *(original source: Maxim)*

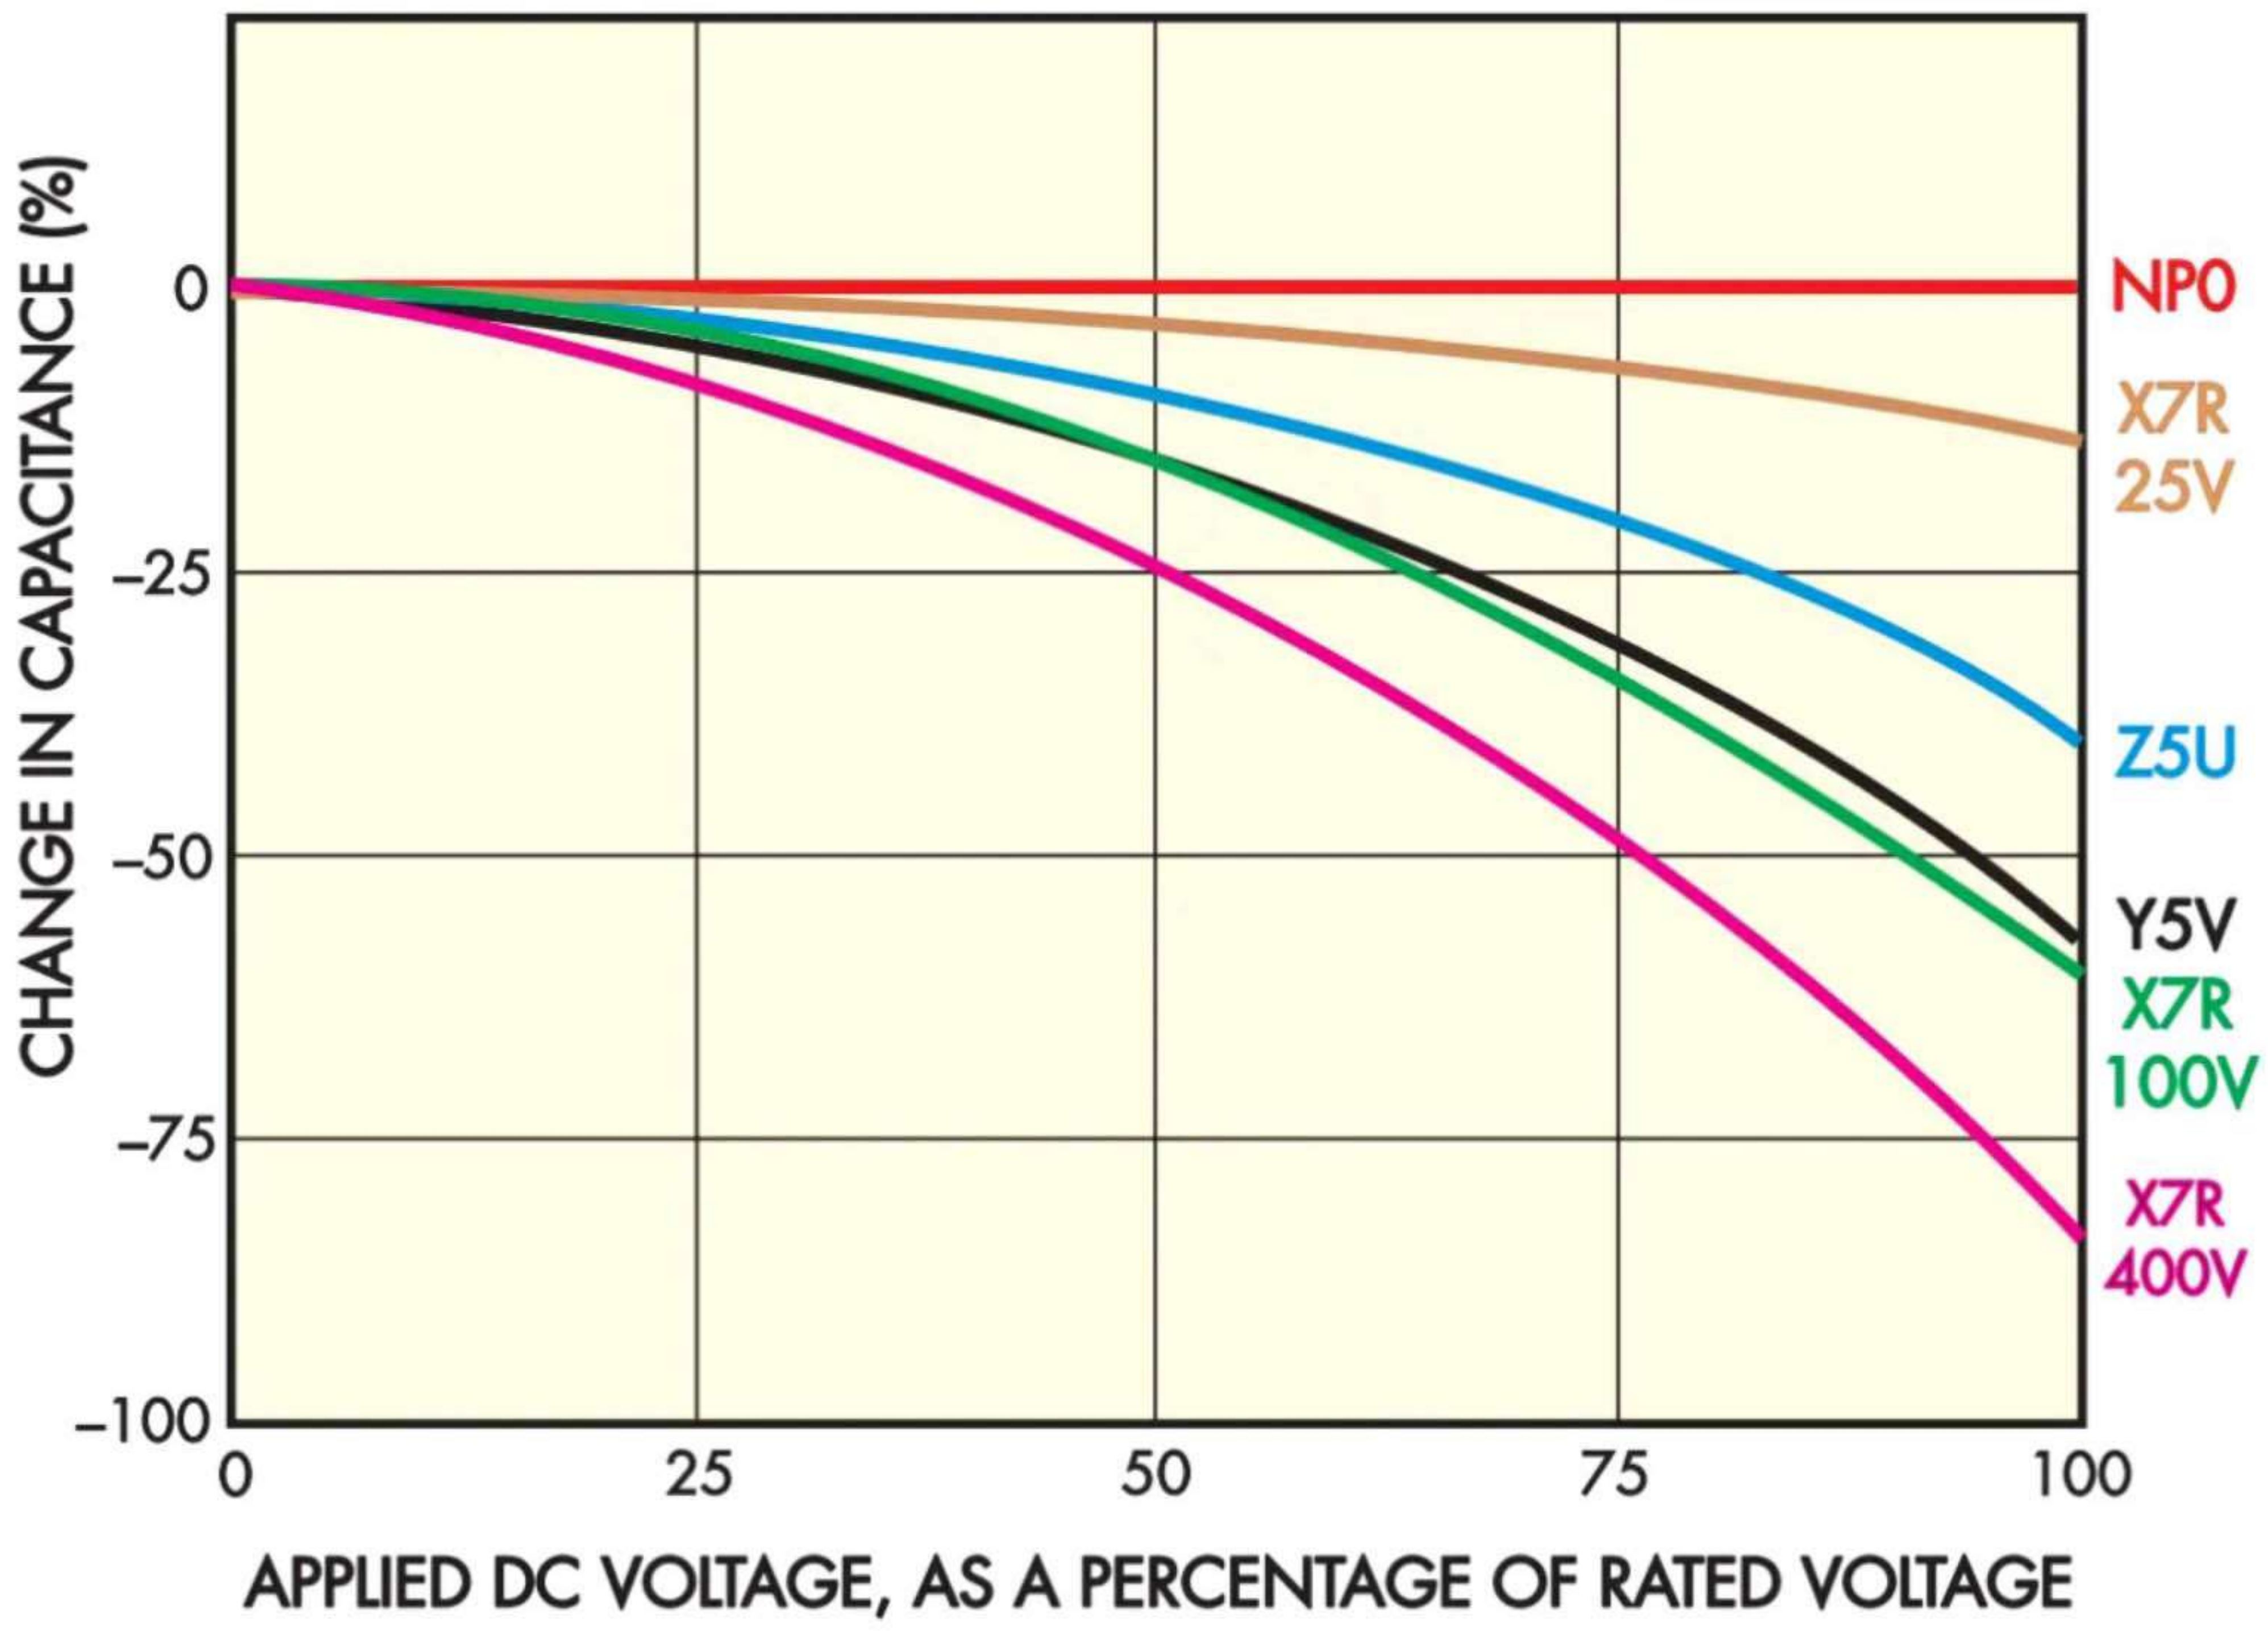

Fig.26: the change in capacitance over voltage for several different ceramic dielectrics. While the 100V and 400V capacitors seem to perform poorly, consider that the X-axis is a percentage of the voltage rating. Due to this effect and the temperature coefficient, Y5V or Z5U capacitors can easily fall below 10% of their rated values, and below even 5% at temperature extremes! (Original source: Wikipedia)

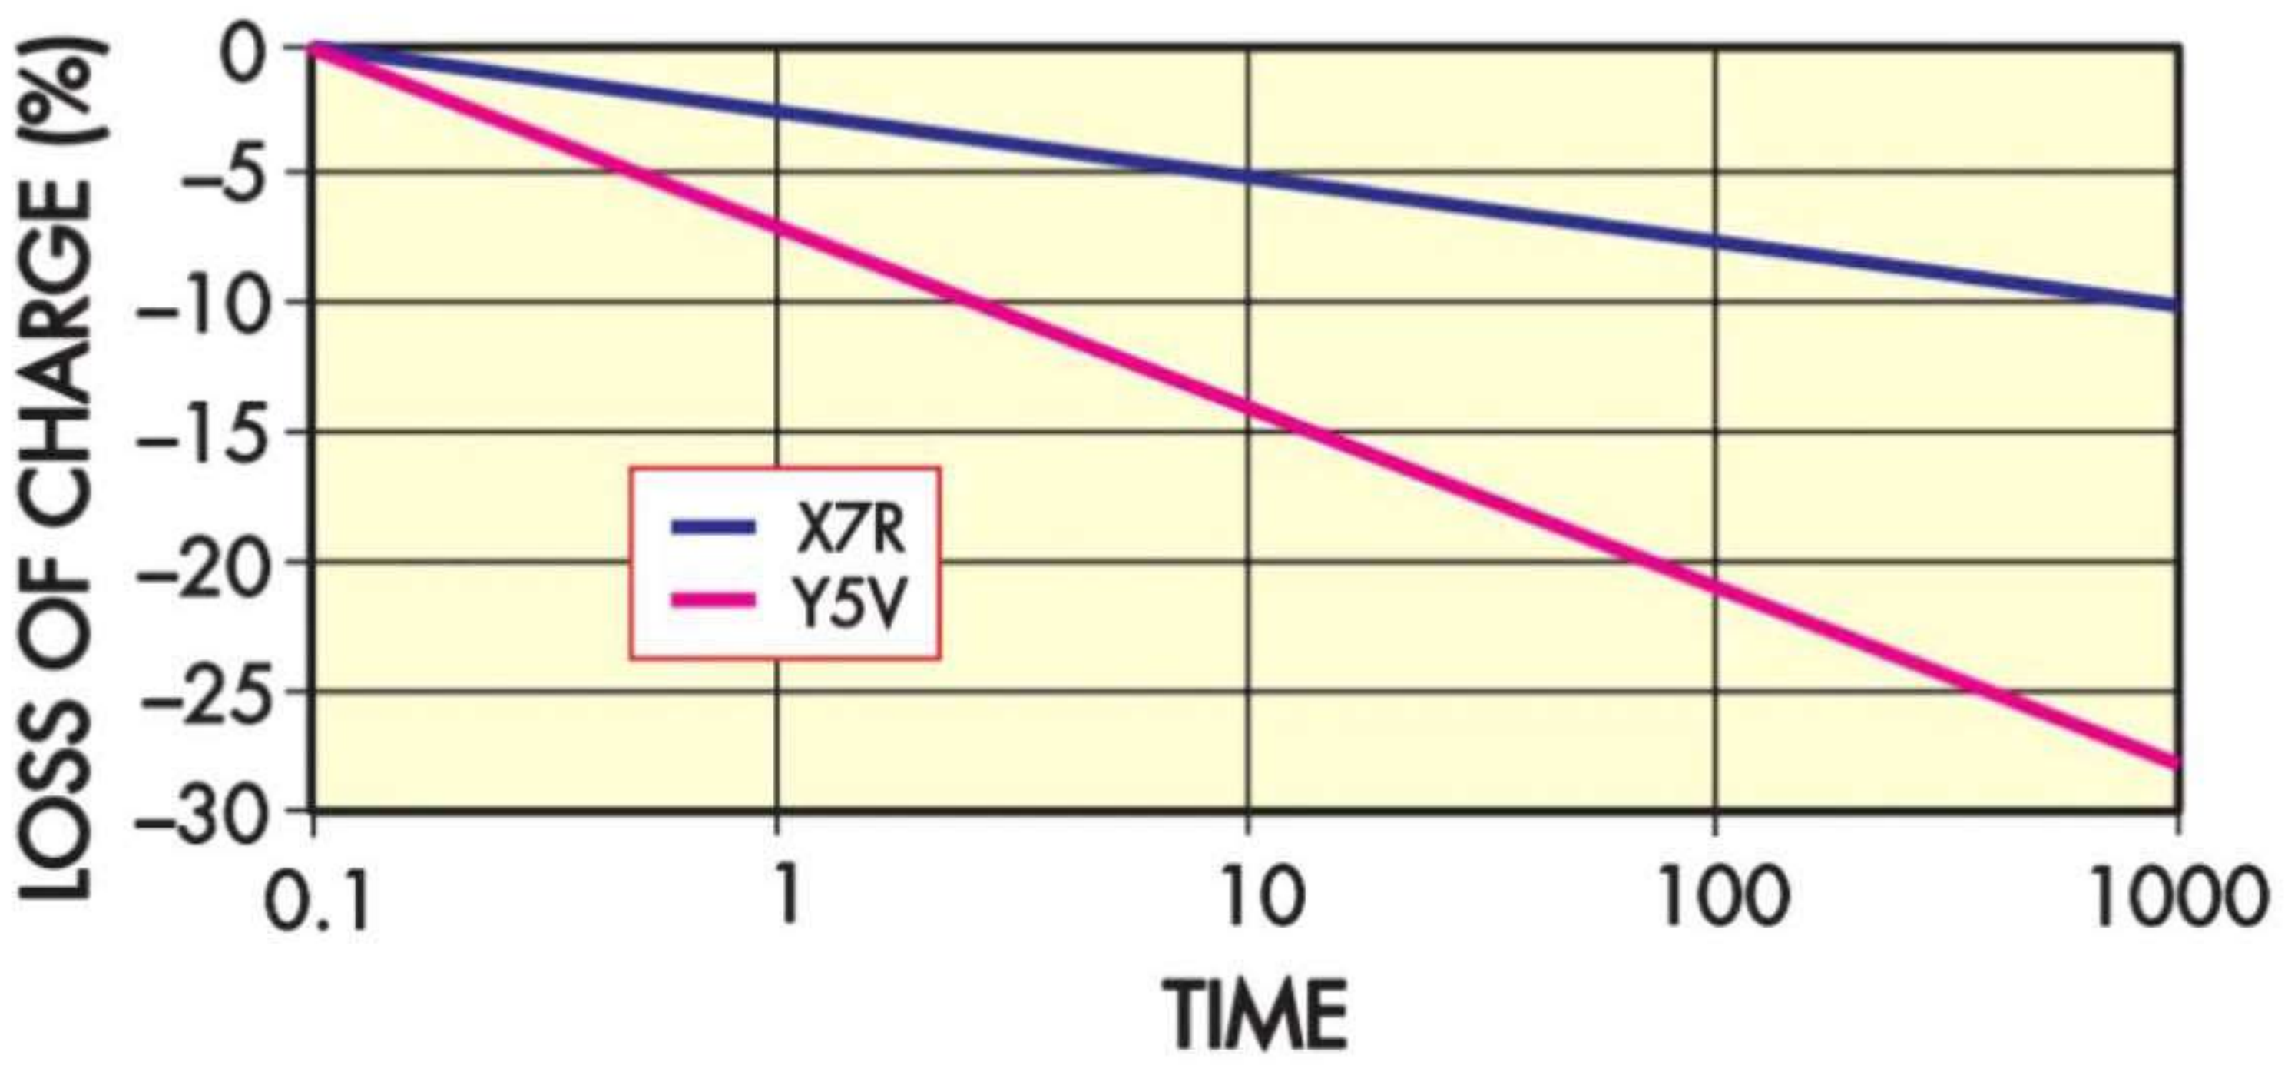

Fig.27: if you need another reason to avoid Y5V ceramics, here is a comparison of the loss in capacitance due to aging with the more robust X7R types. According to the originators of this graph, Johanson Dielectrics, the value of Y5V capacitors drops at roughly three times the rate of a comparable X7R capacitor (original source: Johanson Dielectrics).

Tomorrow's World Today

Techno Talk

Max the Magnificent

I used to avidly watch the weekly *Tomorrow's World* TV program in the 1970s. I remember being blown away by things like pocket calculators and digital watches. I can't imagine what the presenters of that program would have thought about today's technologies!

It's funny the things you remember from decades ago. It's also scary how many things you forget. Then again, maybe I've forgotten how many things I've forgotten – that would be even worse!

Dino dreams

As an example of one of the things that has mysteriously stuck in my mind, there used to be a television series called *Tomorrow's World* when I was a bright-eyed, bushy-tailed young man. First transmitted in 1965 on BBC1, this program introduced the British public to contemporary developments in science and technology.

Some key technologies presented on *Tomorrow's World* that subsequently became commonplace include breathalysers, home computers, audio synthesisers, pocket calculators, digital watches… the list goes on.

I remember watching one program that featured a gigantic robotic dinosaur that looked a bit like a brontosaurus standing in the middle of the TV studio. This sounds much grander than it was. In fact, this 'dinosaur' involved a simplistic static skeleton-like metal framework for the legs and tail. A hydraulic-powered neck was attached to this base. This neck could move side-to-side and up and down.

The 'head' basically involved four tennis-ball-sized microphones mounted in a diamond formation. A computer used a rudimentary form of digital signal processing (DSP) to determine the direction of origin of any sounds based on their time of arrival at the different microphones.

I think this must have taken place close to Christmas circa the mid-1970s because they brought a choir into the studio. Standing in an arc in front of the dinosaur, the choir started singing, quietly at first, then gradually building in volume. They were singing *Silent Night* as I recall, although I may be mistaken.

The dinosaur reacted as soon as the first notes sounded, lowering its head to the height of the singers and extending its neck as far as it could. It's head then traversed back and forth to the left and right as different portions of the choir increased or decreased in volume.

Although this sounds rather pathetic now, it was extraordinary for the time. It inspired a desire in me to work with computers and robots, and it's a memory that will remain with me so long as I manage to maintain a corporeal presence on this plane of existence.

They won't believe you

If you've ever designed a printed circuit board (PCB), you will have heard of the "Gerber format". This open, ASCII, vector format is the de facto standard used by PCB design software to describe the printed circuit board images: copper layers, solder mask, legend, drill data etc.

I never really thought about the origin of this format's name until I became acquainted with David Gerber around 10 years ago. David had just published a book called *The Inventor's Dilemma* about his father, Holocaust survivor Heinz Joseph Gerber.

Heinz was a prodigious inventor. After immigrating to America, he revolutionised antiquated industries and created new ones, like digital plotters. Since there was no infrastructure for this sort of thing, Heinz had to create everything from the ground up, including the Gerber format he used to represent the plots inside the computer.

Compared to Heinz, makers these days have things easy. Rather than being obliged to build everything from scratch, we can take advantage of off-the-shelf platforms like the Arduino or Raspberry Pi.

We have access to small, incredibly accurate, inexpensive microelectromechanical systems (MEMS)-based sensors like three-axis accelerometers and three-axis gyroscopes. We can use open-source languages and file formats to share programs and data with others. The list goes on.

Just the fact that we have immediate access to application notes, data sheets, tutorials, and other documentation via the internet is mind-boggling to people of my generation. Today's younger makers don't realise how lucky they are!

I hail from Yorkshire. I'm reminded of the classic Four Yorkshiremen sketch by Monty Python that ends with "… and if you tell that to the young people today, they won't believe you!" It's true. They don't!

Everyone and everything

Speaking of things that are difficult to believe, implementing things like Wi-Fi, cellular and satellite communications in hobby or professional projects used to be incredibly complex and expensive, but times have changed.

About 10 years ago, a startup called Particle.io introduced their Spark Core Arduino-compatible Wi-Fi-enabled cloud-powered development platform that made creating internet-connected hardware easy and affordable.

A couple of years later, they did something similar with their Electron cellular development kit that came equipped with a SIM card and a global data plan, costing only a few dollars a month.

Just a few months ago, they introduced a dev kit that supports Wi-Fi, Cellular, LoRaWAN and Satellite communications. And, just a few days ago, they… well, we'll save that for a future column.

Did the ground just move?

In my previous *Techno Talk* column, I made mention of the fact that weather forecasts were something of a hit-and-miss affair when I was a kid in the 1960s.

Recently, while watching the news on TV, I saw a report on an earthquake that had just occurred somewhere in Asia. It reminded me of a time in the mid-1980s when I found myself in a hotel room high up in a skyscraper hotel in Tokyo.

I think I was on the 23rd floor. First, the windows started making strange sounds, then the walls started creaking, and the room started shaking, then the whole hotel started swaying. All I can say is that it was an interesting experience.

The reason I mention it is that my wife (Gina the Gorgeous) has iPad apps that allow her to watch graphical depictions of lightning strikes and earthquakes occurring around the world in real time.

I can only imagine what the meteorologists and seismologists back in the 1960s will say when I finally get my time machine working and show them where technology will take us!

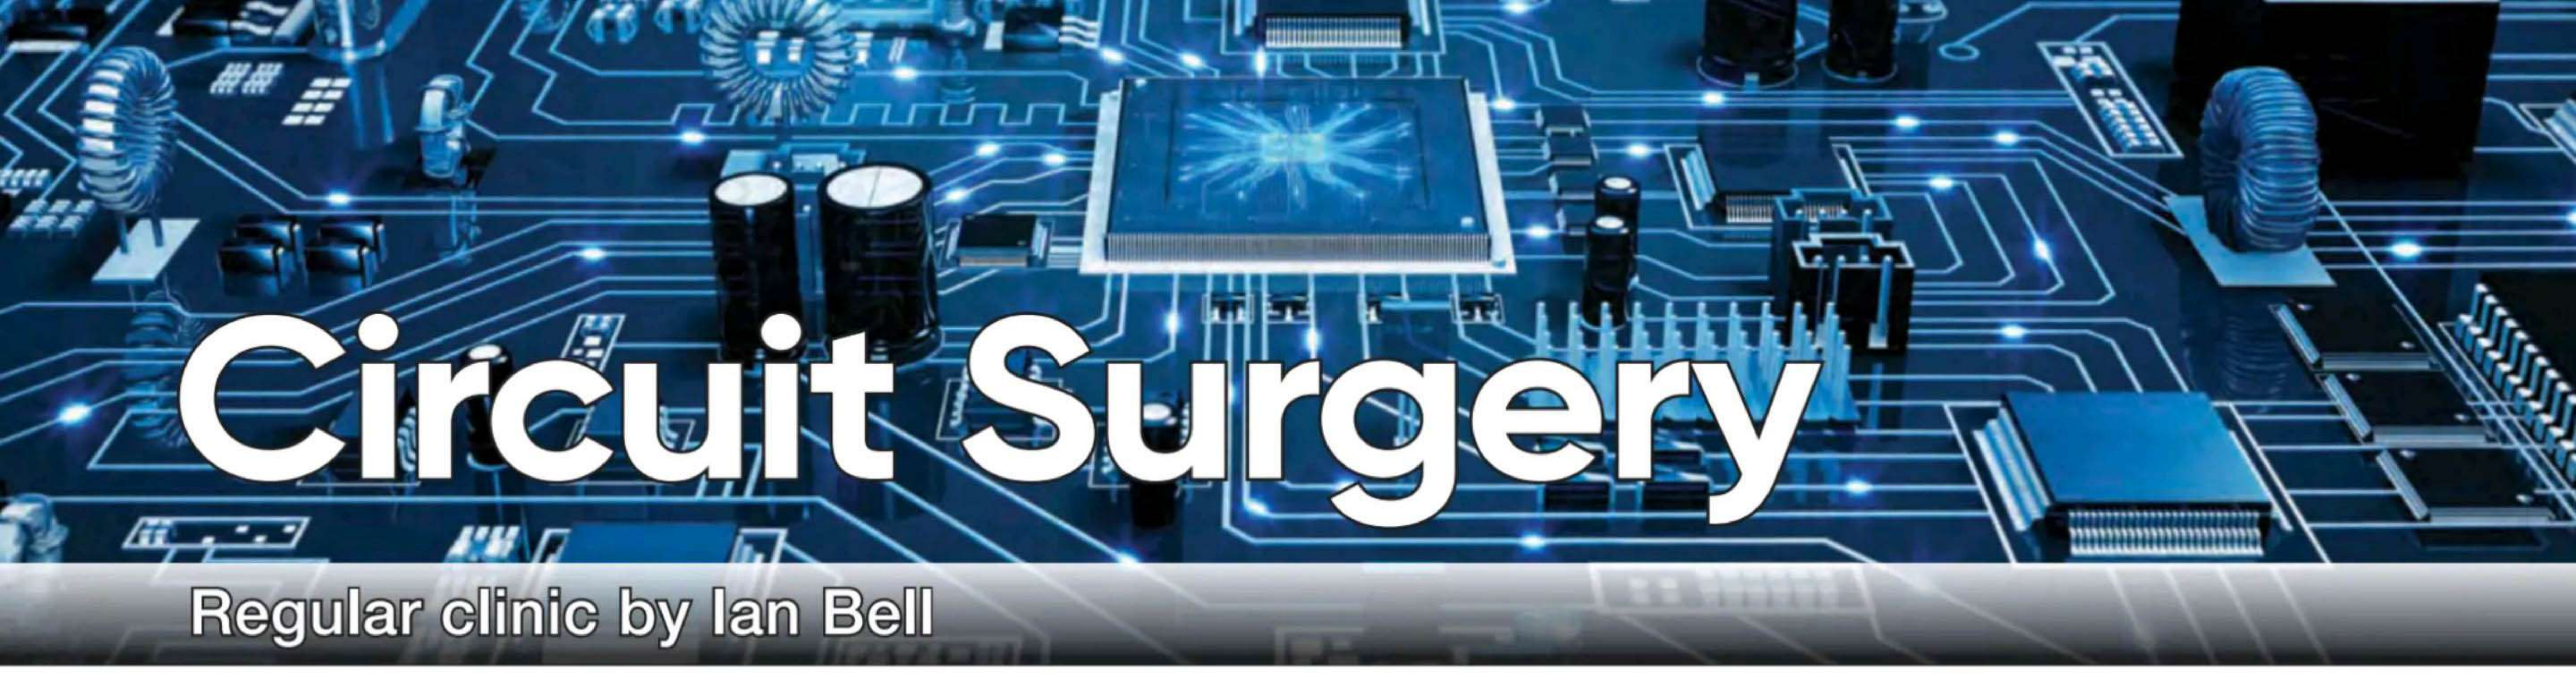

Circuit Surgery

Regular clinic by Ian Bell

Topics in digital signal processing – an introduction to processing digital signals

We are looking at various topics related to digital signal processing (DSP).

DSP covers a wide range of electronics applications where signals are manipulated, analysed, generated, stored or displayed as digital data but originate from and/or are converted to real-world signals for interaction with humans or other parts of the physical world.

Fig.1 shows the key elements of a generic DSP system with a signal path from an analog input via digital processing to an analog output. This does not necessarily represent every DSP system (not all have all the parts shown), but it serves are as reference for the various subsystems we will consider.

In previous articles, we have concentrated on the analog parts of the system (filters and sampling), and on the data converters. This month we will take a look at some basics of digital/sampled signal processing and look at a simple example: a moving average filter.

Sampled signal processing

Consider sampling a continuous signal x_c that is function of time (t). We can write is as $x_c(t)$.

Assume that the sampling period is T. The sampled signal x_s is also a function of time, so we could write $x_s(t)$, but the samples only exist at certain times, specifically integer multiples of T. That is, $t=nT$, where $n=1,2,3$ etc. So we often write a sampled signal as $x_s(nT)$.

Sometimes, this written as $x_s(n\Delta T)$. The Greek letter delta (Δ) emphasises that T represents a time period rather than a continuous time value. Using nT defines a specific sample time, but n alone is sufficient to identify the sample, so we often just use the sample index and write sampled signal as $x(n)$. We can drop the s subscript as the sample-number indexing implies a sampled signal.

In general, we assume that n ranges from $-\infty$ to $+\infty$ (∞ = infinity). This defines a signal which lasts for all time, which is common in theoretical analysis. Of course, real signals have finite duration.

In this discussion, we will only consider systems operating with one sample period, unlike the oversampling DACs discussed last month.

In signal processing, we typically need to consider specific points in time relative to the current time and often define n (or some other symbol) to represent the current sample point of the signals. So, for example, the input to a system at the current time would be $x(n)$ and the current output $y(n)$, as shown in Fig.2 (the Digital Processing part in Fig.1).

The previous input sample is $x(n\text{-}1)$, the input two sample periods before the current time is $x(n\text{-}2)$ and so on. Similarly, the next value of the input, one sampling interval into the future, is $x(n+1)$, followed $x(n+2)$ and so on.

This is illustrated in Fig.3, which shows the samples around the current time in what we assume is a long-duration signal.

The only sample value directly available to the processing system is the current input value, $x(n)$. However, we cannot do much that's useful with just the current sample. We also need access to earlier samples to implement useful functions such as filters.

To make use of earlier sample data, it must be retained in a memory or digital registers in a circuit. The sample number (n, $n\text{-}1$, $n\text{-}2$) then indicates the memory location or specific register holding a particular sample relative to the current input value.

Storing sample data in memory is equivalent to delaying the signal by multiples of the sampling period. Fig.4 shows the signal in Fig.3 delayed by one sample period from the perspective of the same current time.

The value of the delayed signal at the current time is $x(n\text{-}1)$. If we delayed the signal by two sample periods, sample $x(n\text{-}2)$ would be available at the current time. In general, $x(n\text{-}m)$ represents the signal x delayed by m sampling periods.

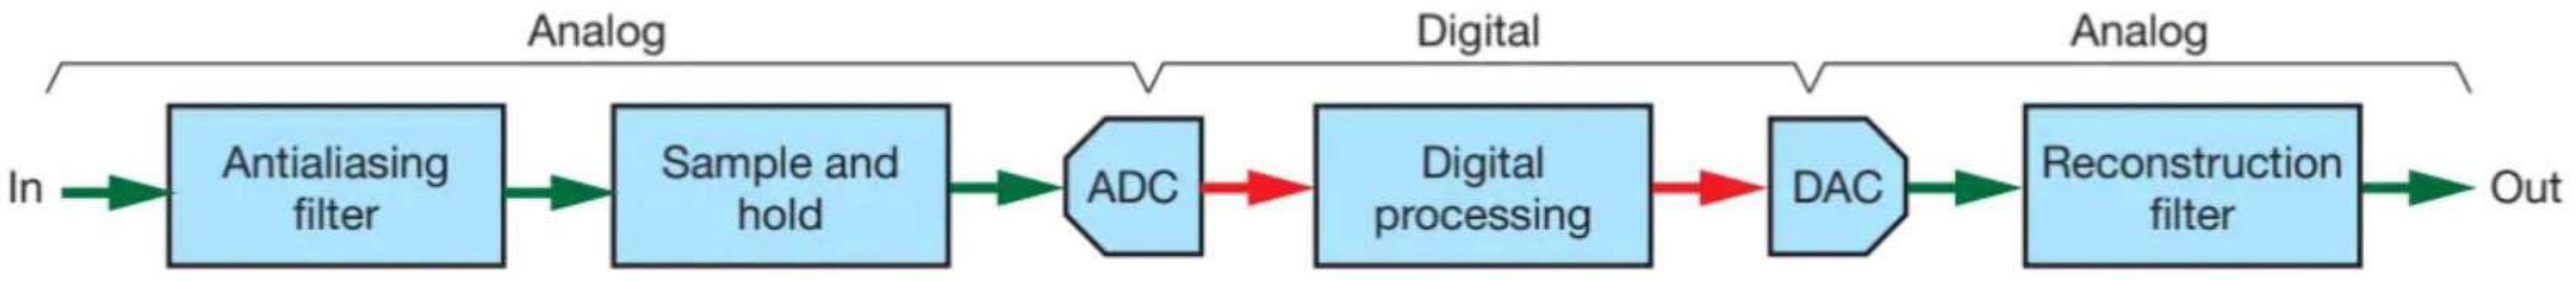

Fig.1: a generic digital signal processing (DSP) system structure.

A signal processing system may also need to use previous output values in the calculations it performs. Memory or registers can be used to store these values in the same way as input values.

Fig.5 shows a generalised digital signal processing system; the input and output are the signals x and y, respectively, at current sample point

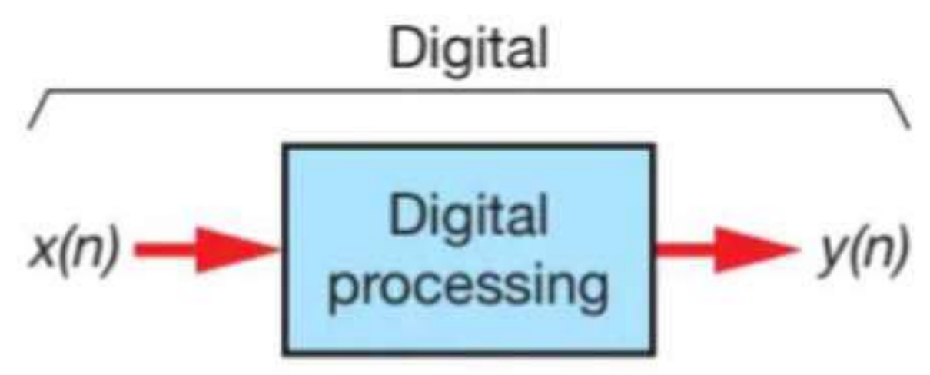

Fig.2: a digital sample processing scheme.

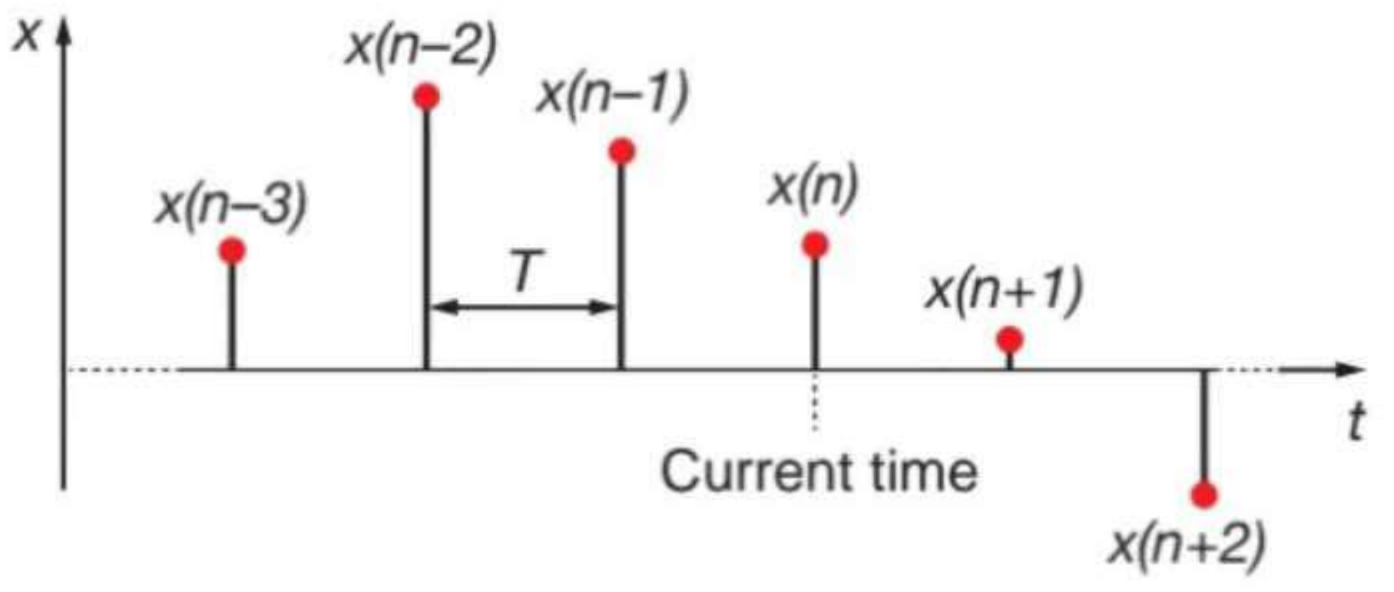

Fig.3: part of a sampled waveform.

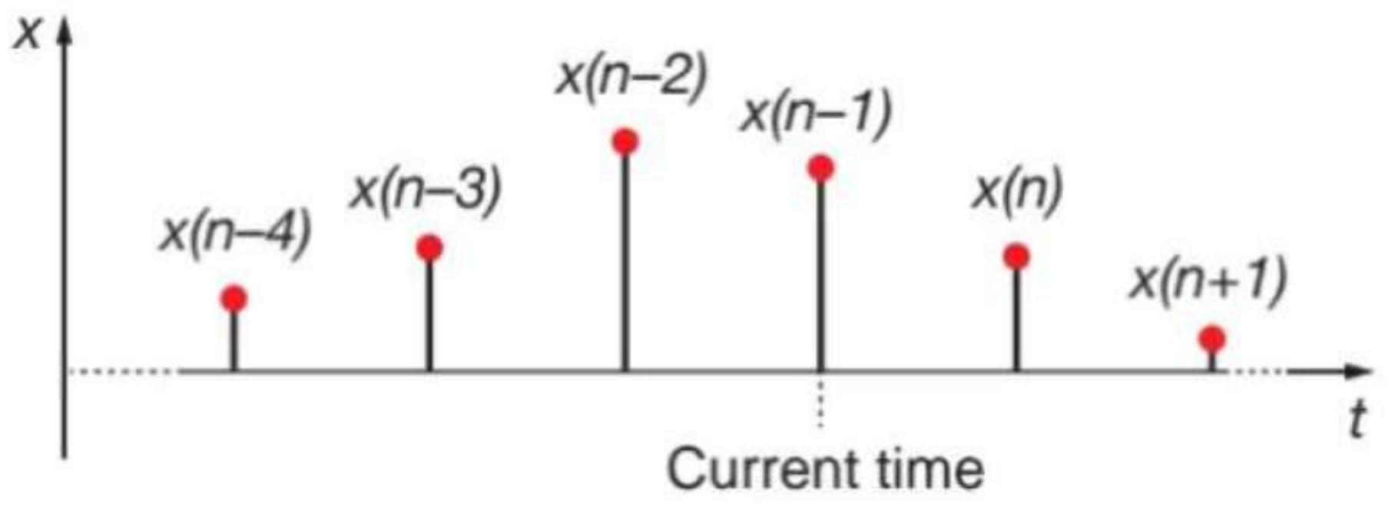

Fig.4: like Fig.3 but delayed by one sample.

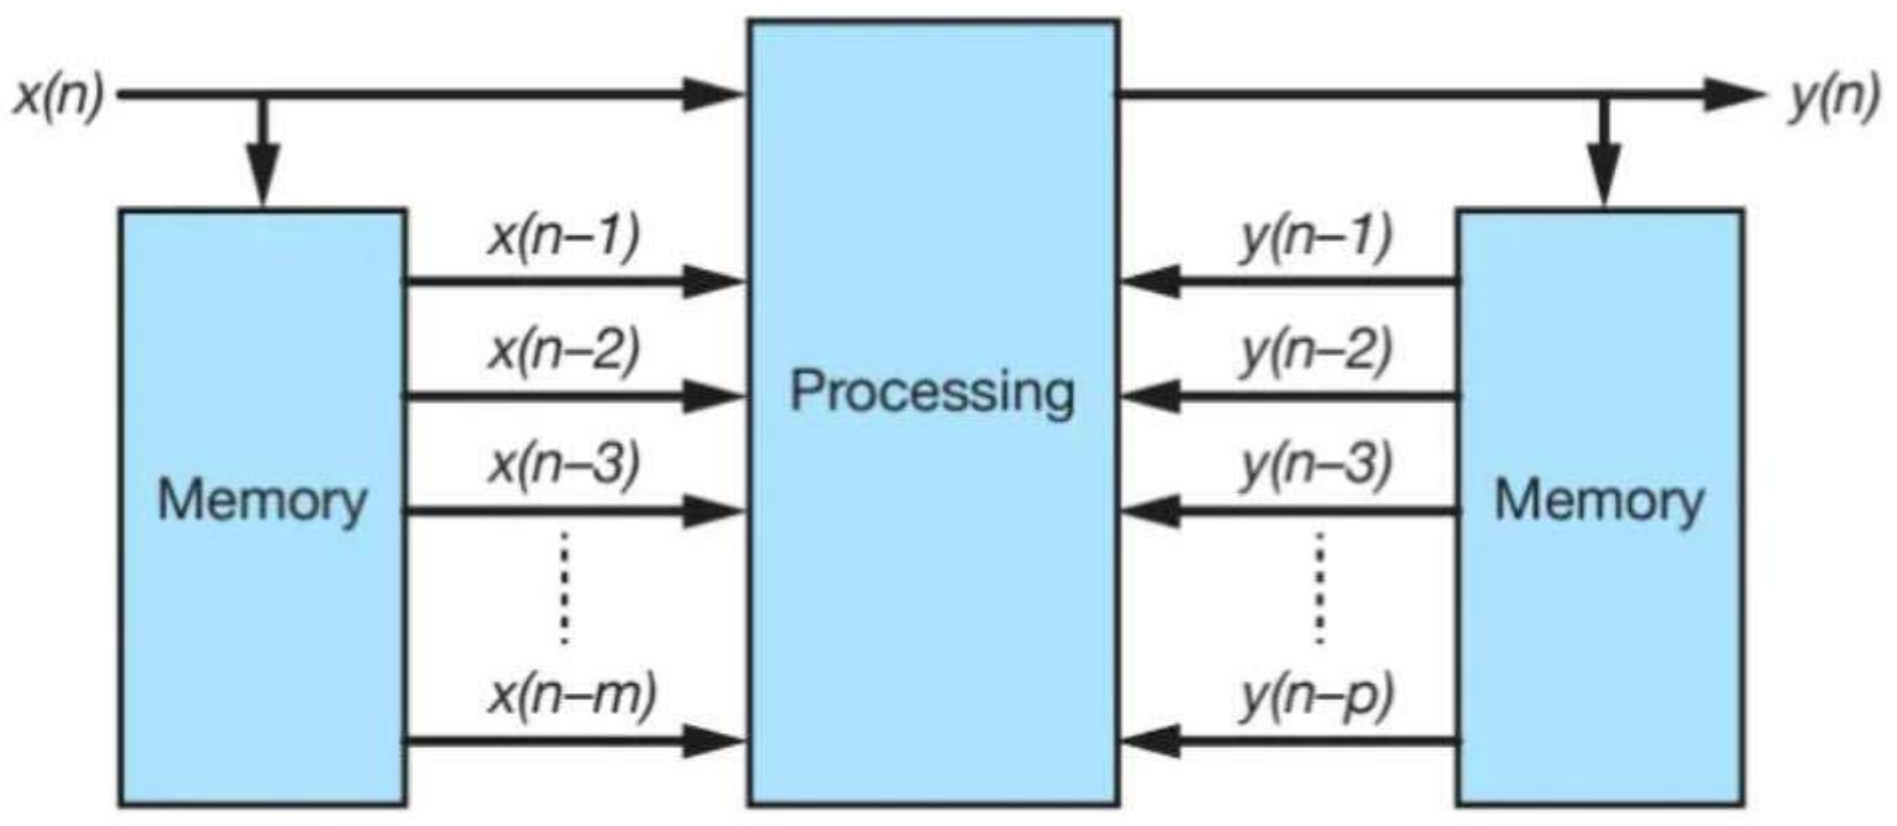

Fig.5: a general digital processing system.

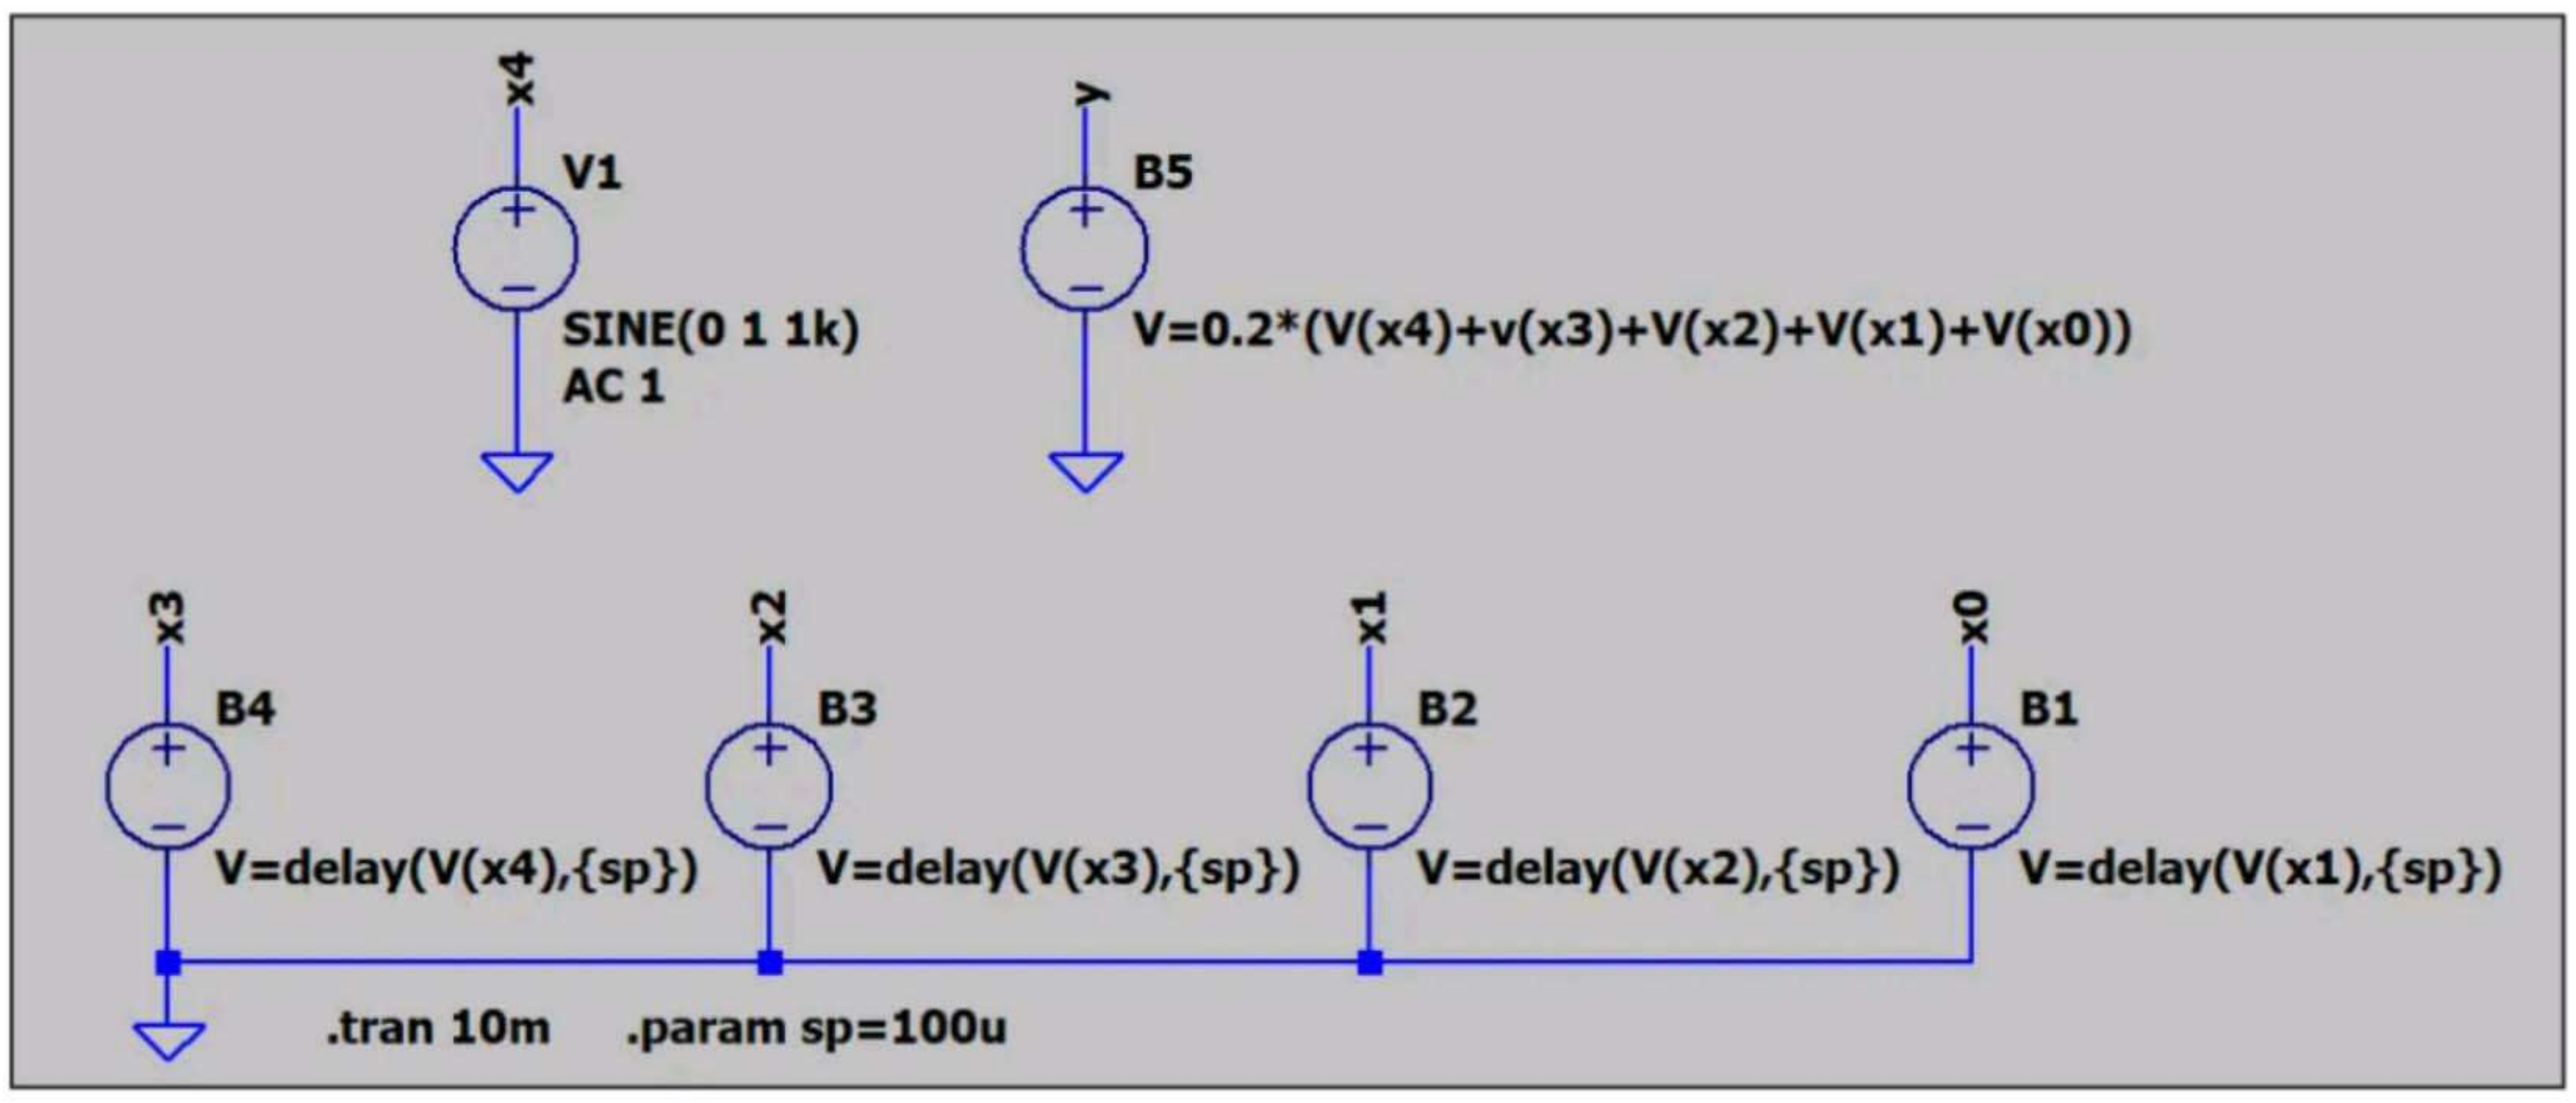

Fig.6: this shows how I implemented signal delays in LTspice.

n [that is, *x*(*n*) and *y*(*n*)]. A total of *m* previous inputs [*x*(*n*-1) to *x*(*n*-*m*)] and *p* previous outputs [*y*(*n*-1) to *y*(*n*-*p*)] are stored in memory and used by the signal processing.

Real systems vary widely in the number of previous input values used, and many do not use any previous output values.

It is not uncommon for the fundamental theory related to signal processing to involve future input values, denoted [*x*(*n*+1), *x*(*n*+2), ...]. Systems defined in terms of future inputs are called non-causal systems; obviously, they cannot be implemented for real-time input signals.

However, theoretically developed non-causal systems can be modified to make them causal; for example, by time-shifting the calculation so that only current and past values need to be used. If the signal being processed is already stored in its entirety (eg, digital audio on a hard drive), then with respect to any given data point, the 'future' values are available.

A moving average filter

Possibly the simplest digital signal processing function is the moving average filter – we will use it as an example for that reason. It is worth quickly investigating its properties first.

Moving averages are widely used in a variety of fields, not just electronics. If you do an internet search on "moving average", the top results are likely to mainly relate to finance and economics. Moving averages are useful because they provide a simple and clearly defined smoothing function for data. Plotting a moving average makes it easier to visualise trends.

Moving averages can also help to remove known periodic fluctuations. An example of this is the use of seven-day moving averages for the number of cases during the COVID pandemic. Test labs may not be operating every day, which creates peaks and troughs over the course of a week. The seven-day moving average removes this effect, as well as smoothing out the natural random variations in case numbers.

In electronics, moving average filters are very effective at removing random noise from signals. They are particularly useful when the signals have fast step-changes in level (eg, pulse-based signals), where the timing and shape of steps needs to be preserved as much as possible through the filtering process. They are not very useful for frequency separation (rejecting some frequency ranges while passing others).

As its name suggests, a moving average filter finds the average (mean) of a range of sample values from the input to obtain the output. This is achieved by summing the samples and dividing by the number of samples. For example, for five samples, the moving average processing function is:

$$y(n)=\frac{x(n-4)+x(n-3)+x(n-2)+x(n-1)+x(n)}{5}$$

The equation relates to Fig.5, in which *m*=4 and in which no previous output values (*y*) are stored and used – only the input memory is present. This is a causal version that can be implemented in real time.

We could also take the average of the current sample, two past and two future ones. This average is centred on the current sample, but is non-causal. However, this may be useful in data analysis if the specific sample number a result is associated with is important (ie, it avoids time-shifting large changes).

An LTspice implementation

In Circuit Surgery, we often simulate circuits using LTspice, so it is worth checking if we could do this for a moving average filter. It would be useful to see the effect on a noisy waveform, and to obtain frequency response curves.

This is perhaps not an obvious thing to use LTspice for, but it is worth trying, and doing so also demonstrates some useful features of the simulator.

We can use a behavioural source to implement the averaging maths if we can produce signals to represent the stored previous waveform values. As discussed above, storing previous samples is equivalent to delaying the sampled signal by an integer multiple of the sampling period.

LTspice behavioural sources support a delay function: `delay(x,t)`, which delays signal `x` by time `t`, enabling us to implement the memory function. An example of this is shown in Fig.6.

In the five-sample moving average equation above, if *n*=4, the terms in the equation are *x*(0), *x*(1), *x*(2), *x*(3) and *x*(4), where *x*(0) is the oldest sample (after 4 sample period delays); the current input signal is *x*(4). These index values are used to name the signals in the LTspice schematic in Fig.6 (nets x0 to x4). Behavioural sources B1 to B4 implement the delays.

Each delay is the same; this value is defined as a parameter, `sp` (for sampling period) using the `.param` directive. This means we only have to change this one value to change the delays of all the behavioural sources.

LTspice parameters are used by putting the parameter name (or an expression using them) inside curly braces (eg, `{sp}`). This is used in the equations for the behavioural sources, for example,

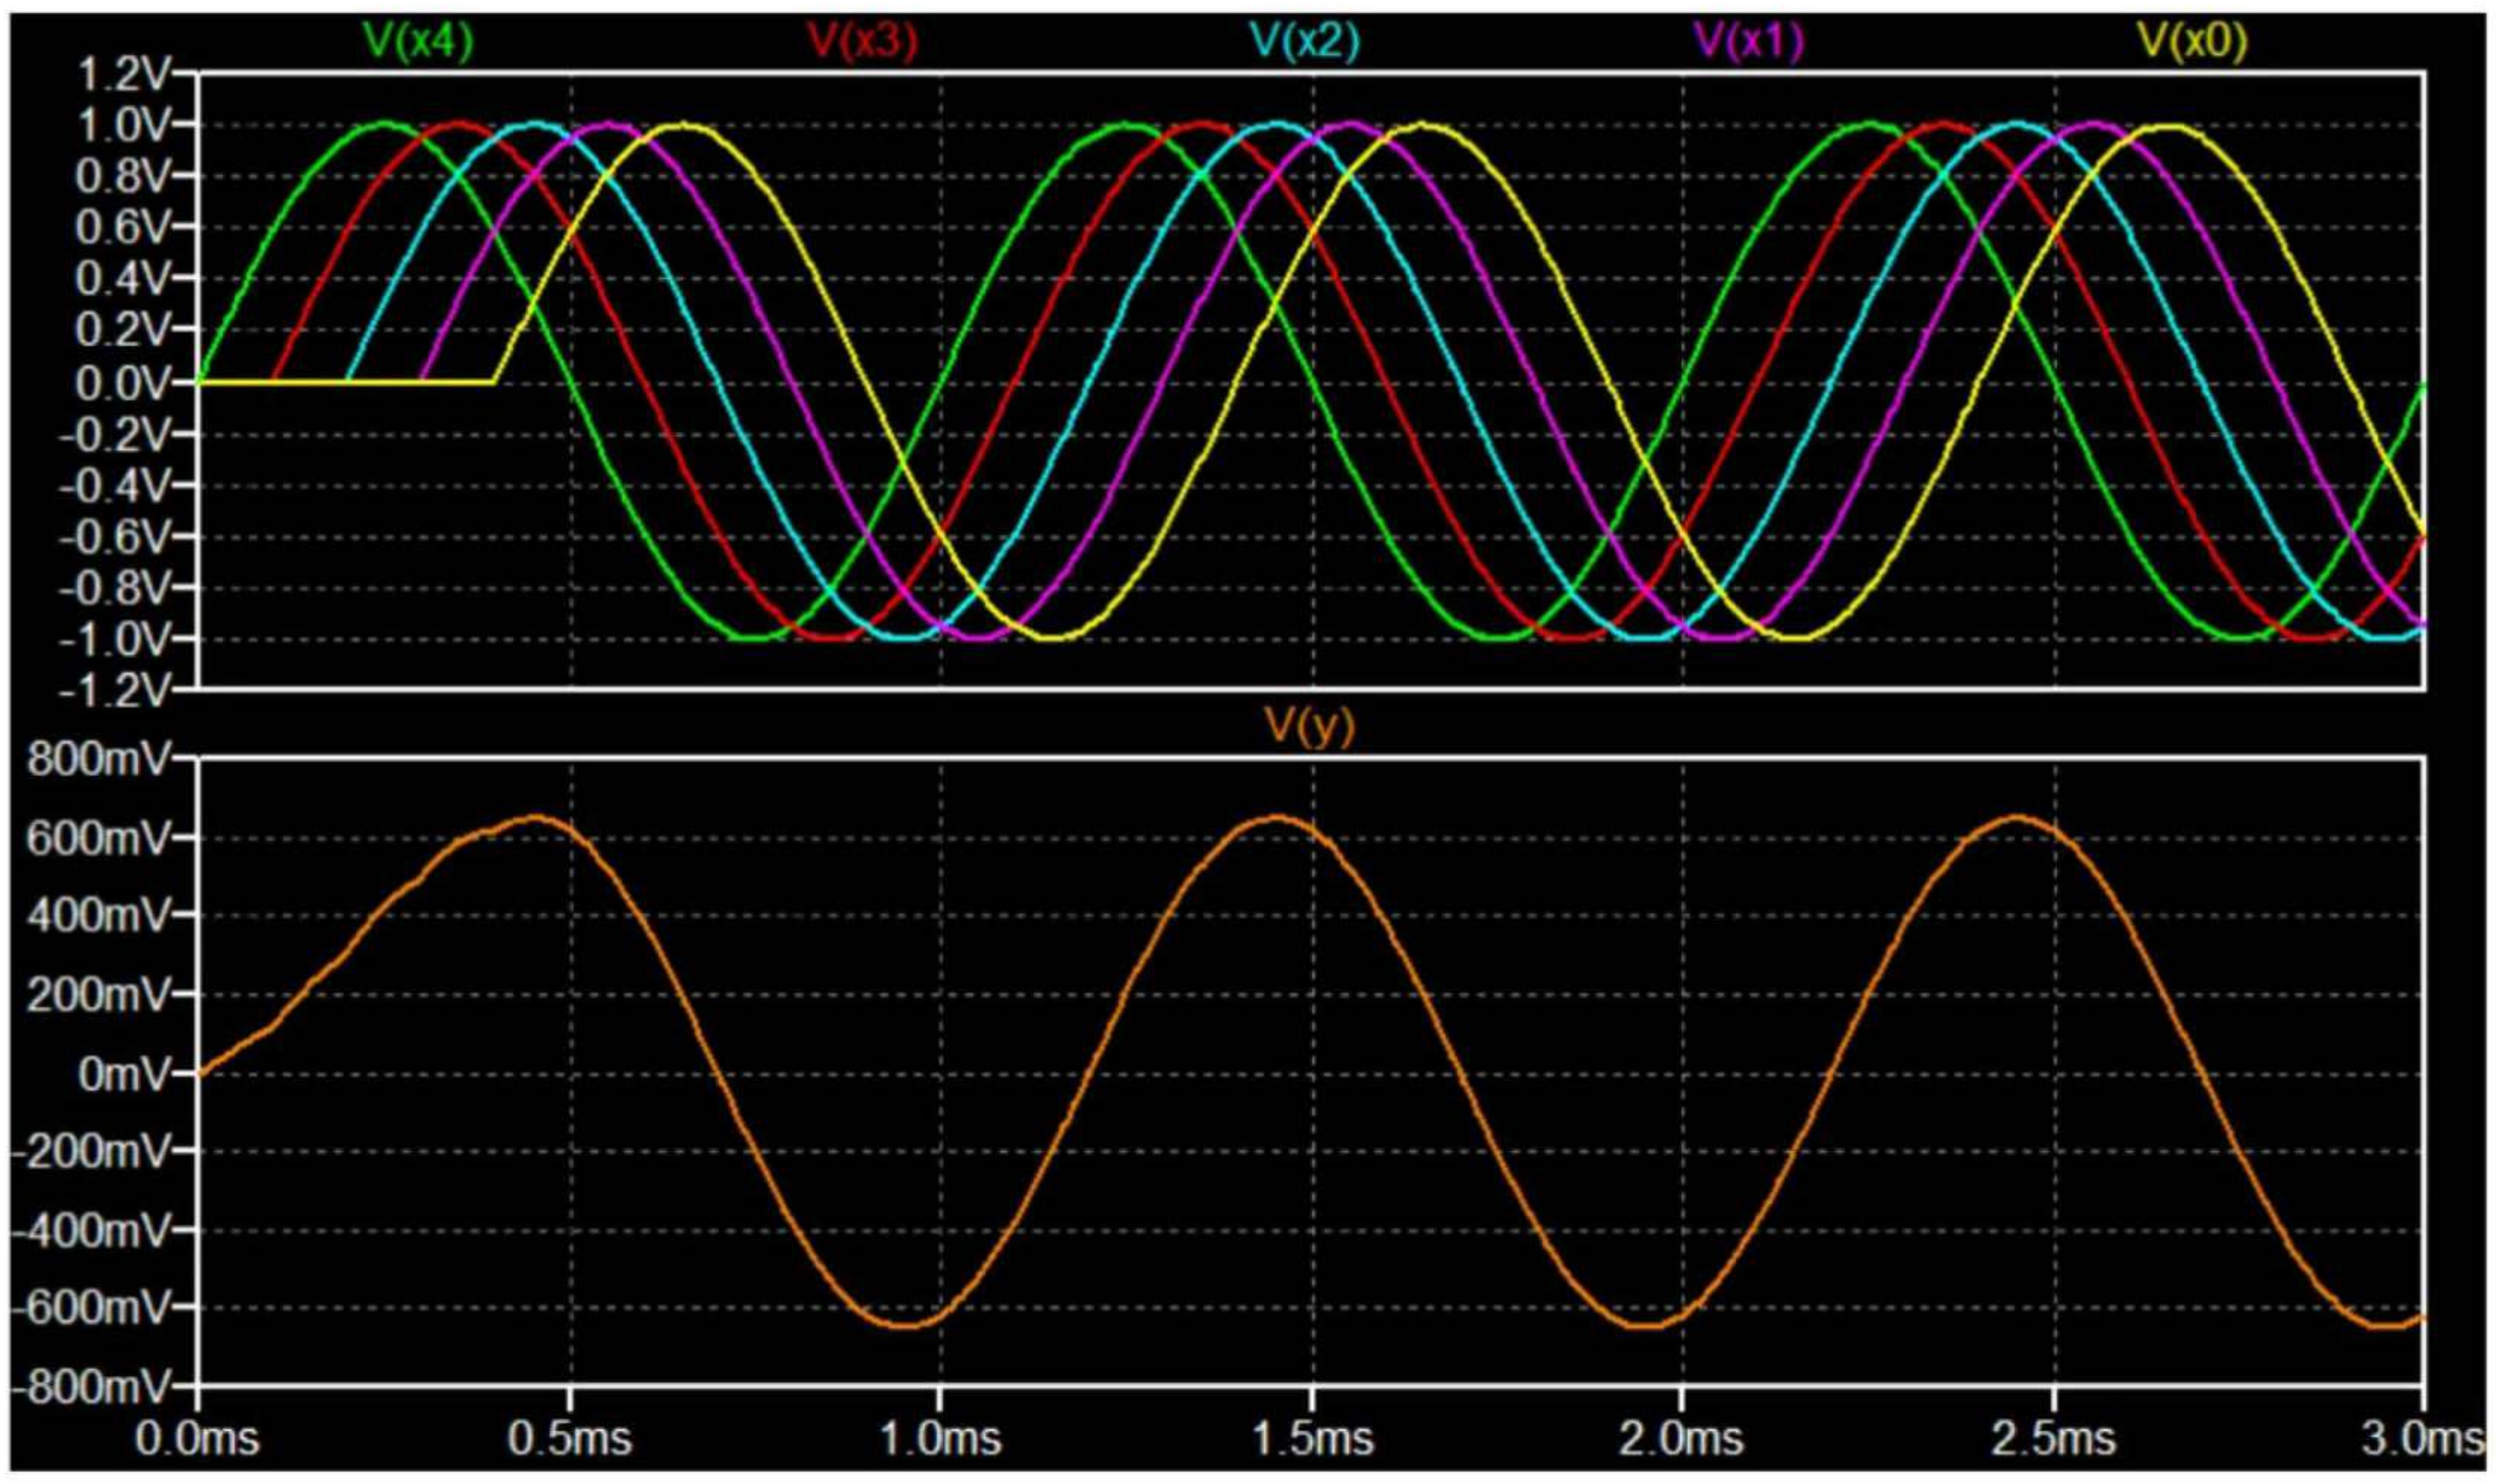

Fig.7: the results of running the simulation on the schematic shown in Fig.6.

B4 generates the voltage on net **x3** as a copy of the voltage on net **x4**, delayed by **sp** using **V=delay(V(x4),{sp})**.

In the Fig.6 circuit, behavioural source B5 creates the output signal **y** by implementing the five-sample averaging function: **V=0.2*(V(x4)+v(x3)+V(x2)+V(x1)+V(x0))**. The input (**x4**) is a sinewave generated by voltage source V1. The results are shown in Fig.7, which shows the evenly-spaced delays on the waveforms.

The output settles to a sinewave because the averaging function is performing a linear operation on its sinewave inputs. However, the waveform is initially not sinusoidal because, at the start, the averaging uses the initial values of the delayed waveforms, which are zero rather than sinusoidal.

It is common in digital signal processing for the outputs to not be meaningful until all the sample memories have values relating to the actual input signal in them. Usually, this is not a problem, and designs are based on continuous processing of a long-lasting signals.

The circuit in Fig.6 is missing the filtering and sampling parts of a typical DSP system, as per Fig.1. These can be added using the "sample" device we discussed previously to produce a sampled waveform, along with behavioural filter devices to provide the anti-aliasing and reconstruction filters. These were discussed in previous articles in this DSP series.

Fig.8 shows the components added to the schematic in Fig.6 to implement sampling and filtering. Sinewave source V1 is now connected to the net **in**, with net **x4** (current sample) now coming from the output of the sample-and-hold block. All the behavioural sources from Fig.6 (B1 to B5) are also present in the simulation schematic, but are not shown in Fig.8.

In the Fig.8 circuit, the sampler (A1) and filter devices (U1 to U4) use the sampling period parameter value to control their operation. The sampling clock period produced by source V2 is equal to **sp** and the cutoff frequency of the filters is set to the Nyquist frequency (half the sampling frequency) at **1/(2*sp)**.

The results from simulating Fig.8 are shown in Fig.9. The upper plot pane shows the sampled sinewaves with different delays. If you look at the green waveform (**x4**, the current input to the processing) at a given sample point, you can see the same sampled level appearing on the other waveforms in sequence (red, cyan, magenta, yellow) as the value moves through the delay chain.

The lower pane shows the sampled output (**y**) and the final output after passing through the reconstruction filter (**y_filt**). The filtered output is very similar to the output of the non-sampled version of the system in Fig.7.

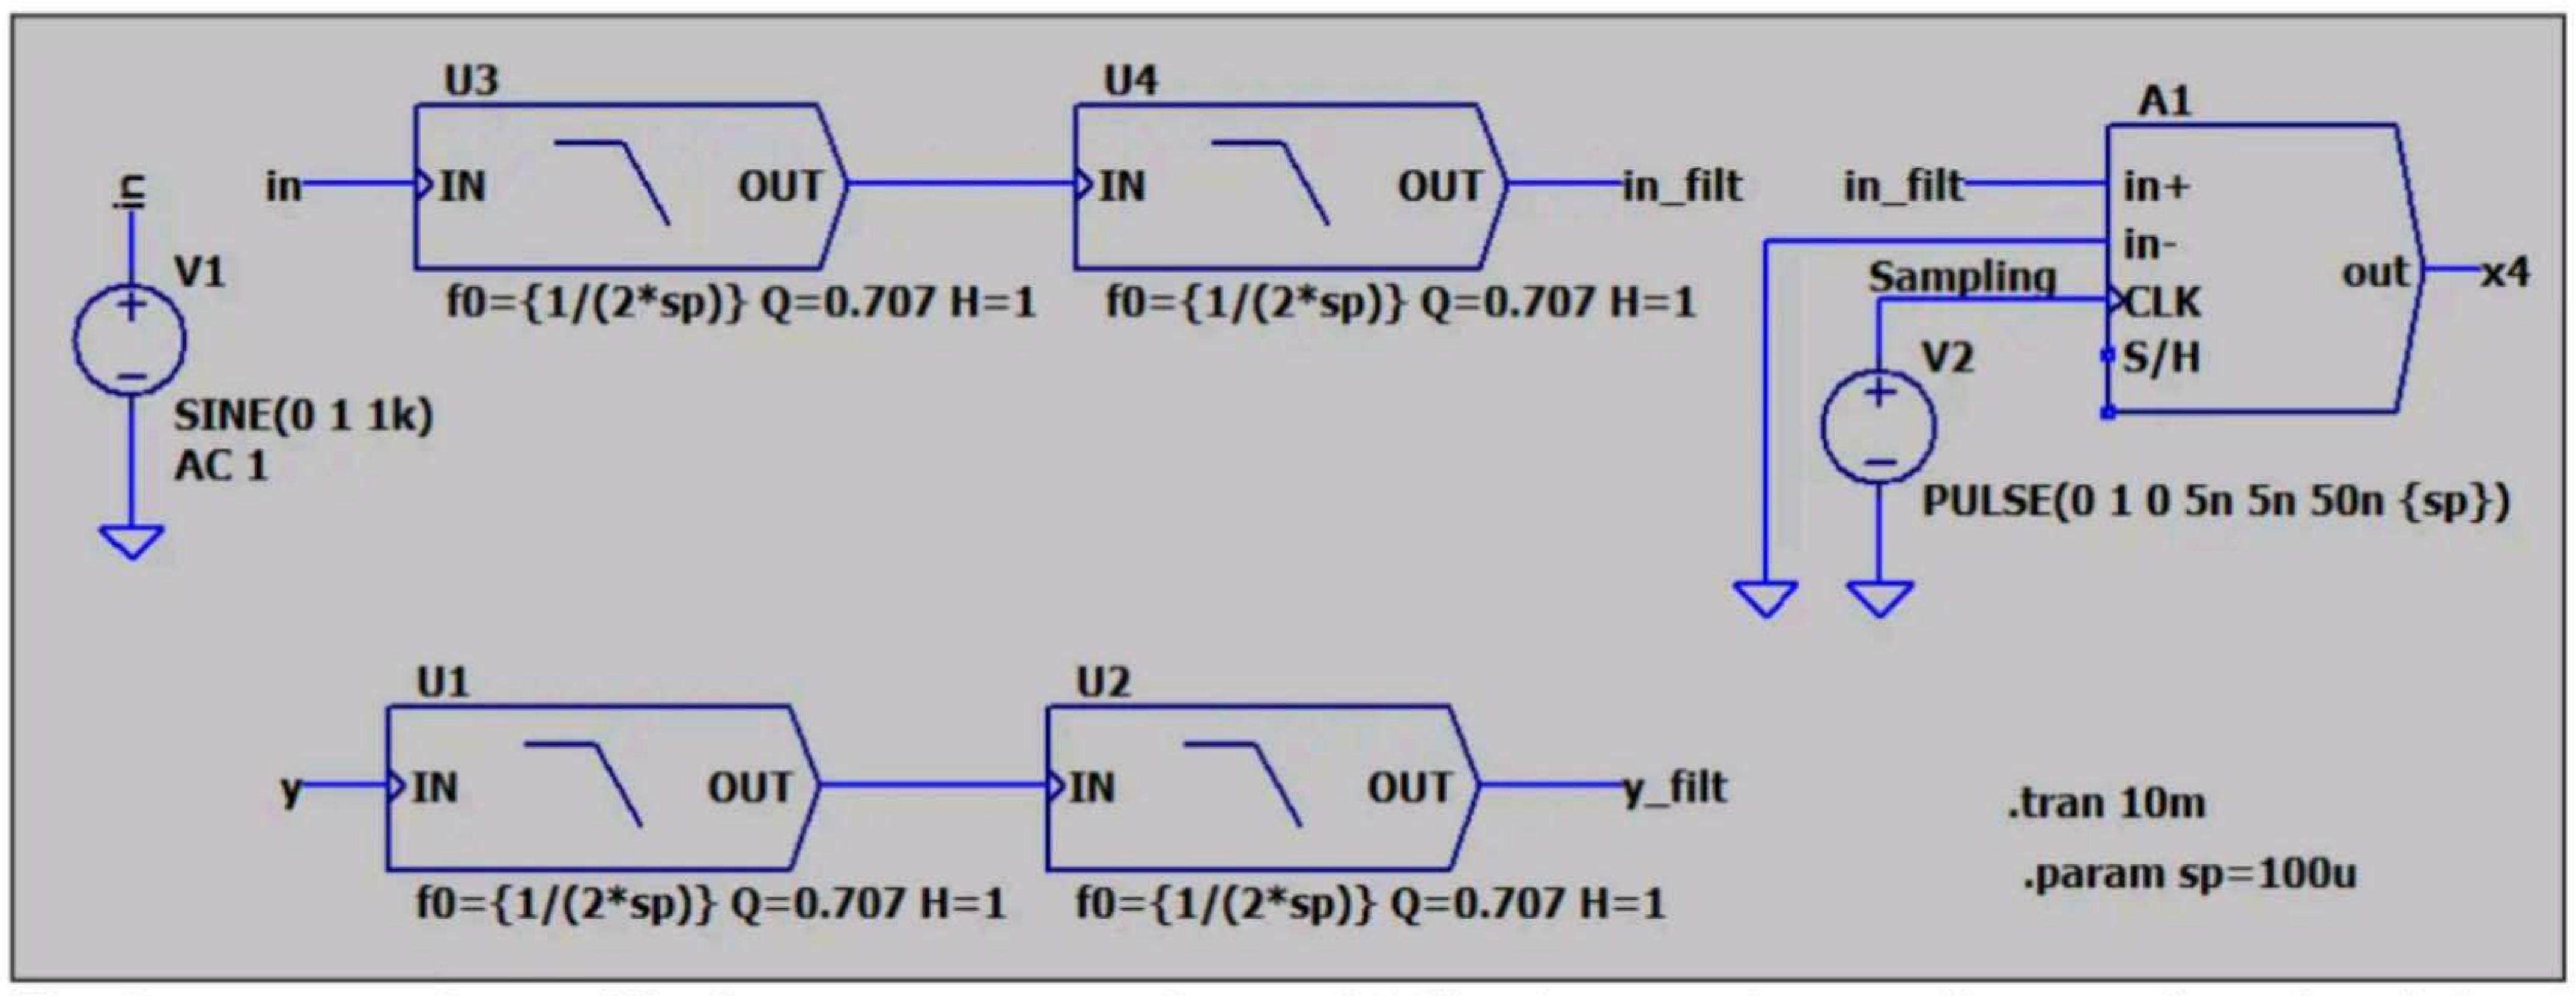

Fig.8: new and modified components that add filtering and sampling to the circuit in figure 6. B1 to B5 are also present but not shown here

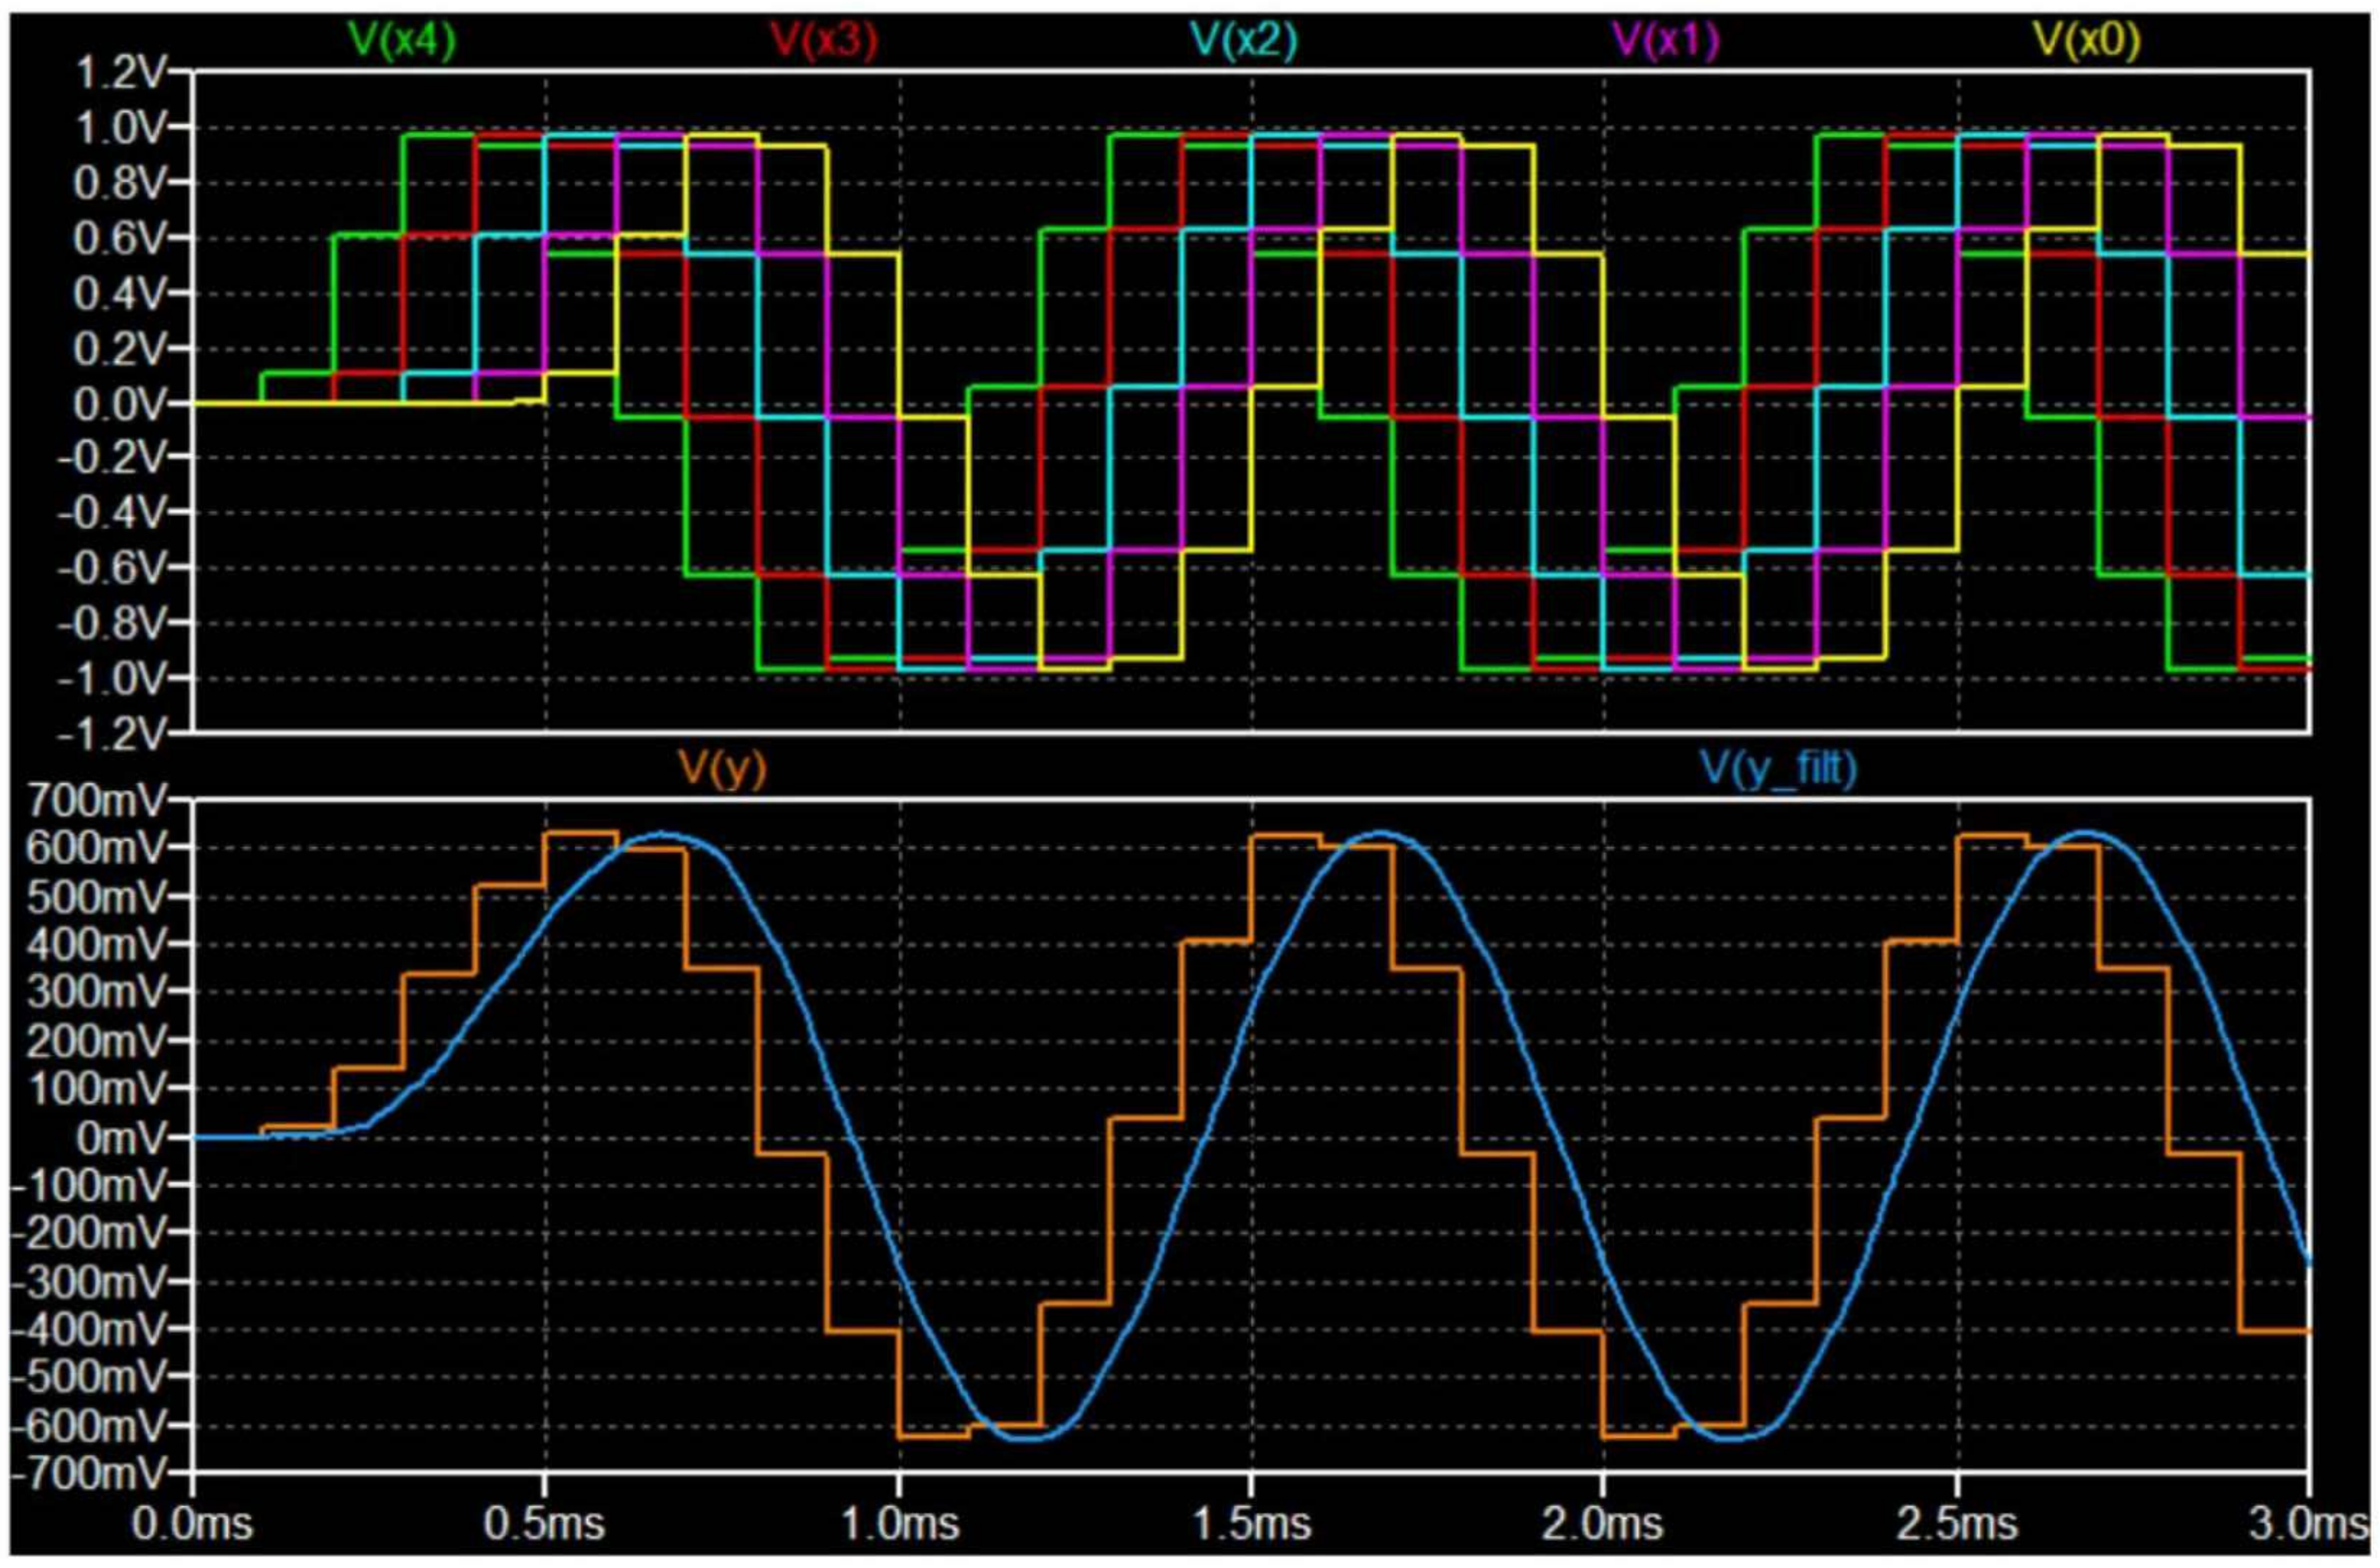

Fig.9: the results from the schematic shown in FIg.8.

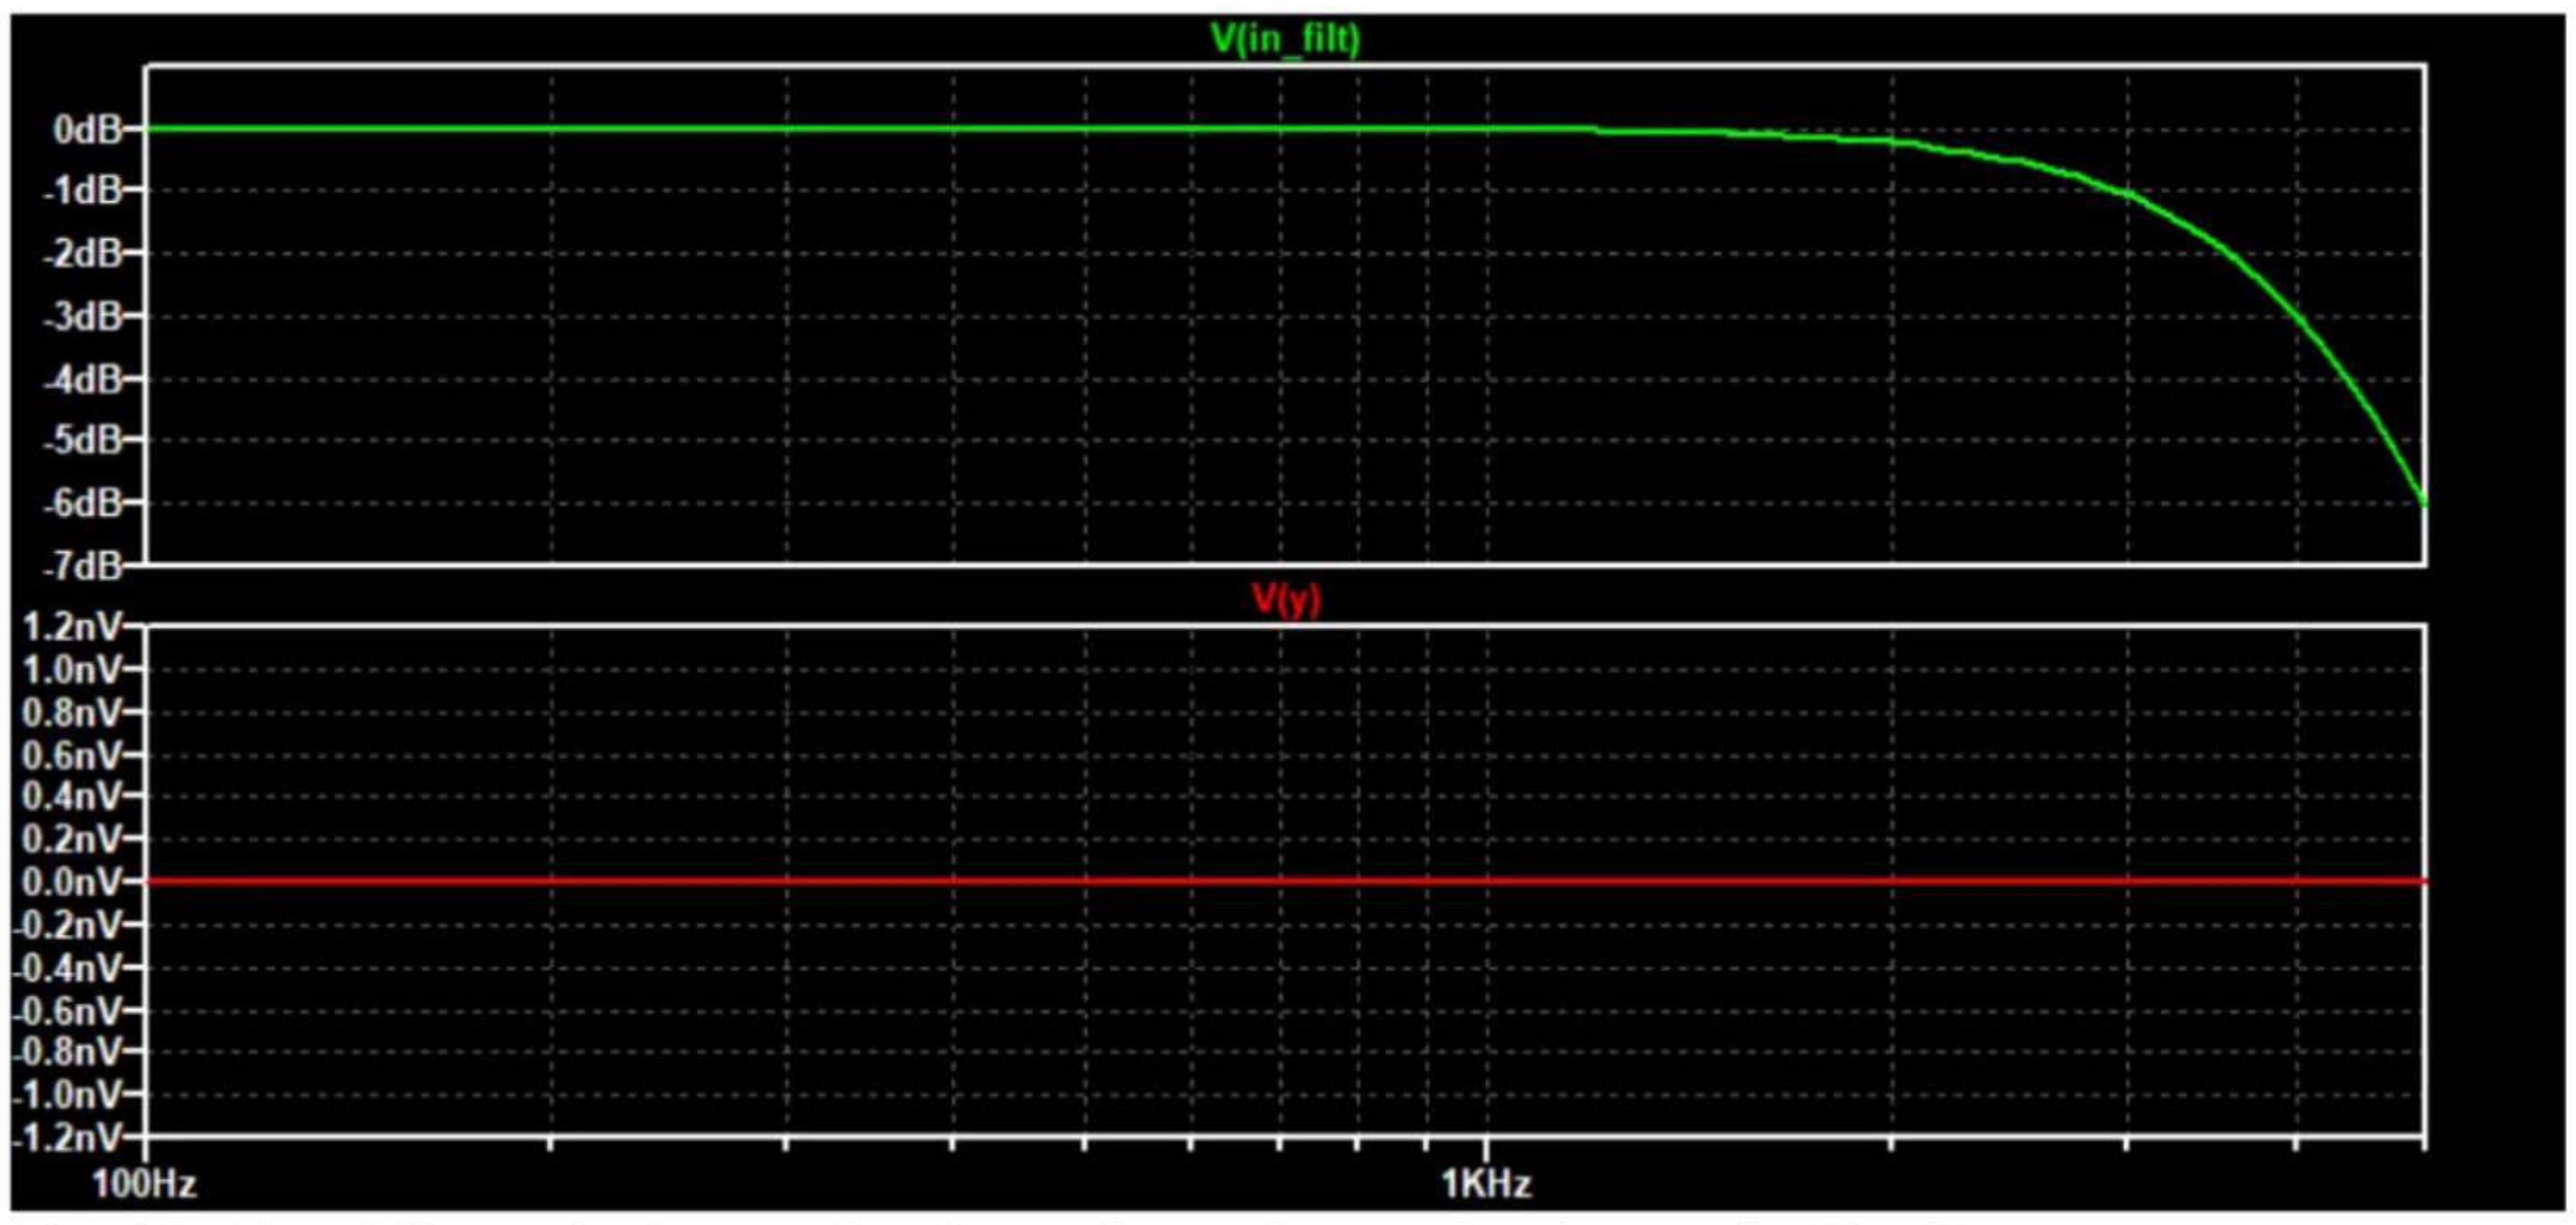

Fig.10: the AC analysis results from the schematic shown in Fig.8.

AC analysis failures

Given that the moving average function acts as a filter, we might want to look at its frequency response. Typically, we would use AC analysis in LTspice to achieve this.

The V1 source in Figs.6 & 8 is configured ready for this, with an AC amplitude of 1V (**AC 1**). All we need to do is configure the AC analysis, for example, using **.ac dec 50 100 5k** to plot the response from 100Hz to the Nyquist frequency of 5kHz. The sampling period (**sp**) is 100µs, so the sampling frequency is 1 ÷ 100µs = 10kHz.

The results are shown in Fig.10. The upper trace is for **in_filt** and therefore shows the response of the anti-aliasing filter. This is as expected from the setup of the behavioural filters, and shows very little attention at the 1kHz frequency of the sinewave from V1 used as the input of the simulation in Figs.7 & 8.

The lower trace, the output of the signal processing (**y**), is zero at all frequencies. The AC analysis is not working because it is not supported by the sample-and-hold device. The switched nature of the sampling means that we cannot use basic AC analysis to determine the frequency response.

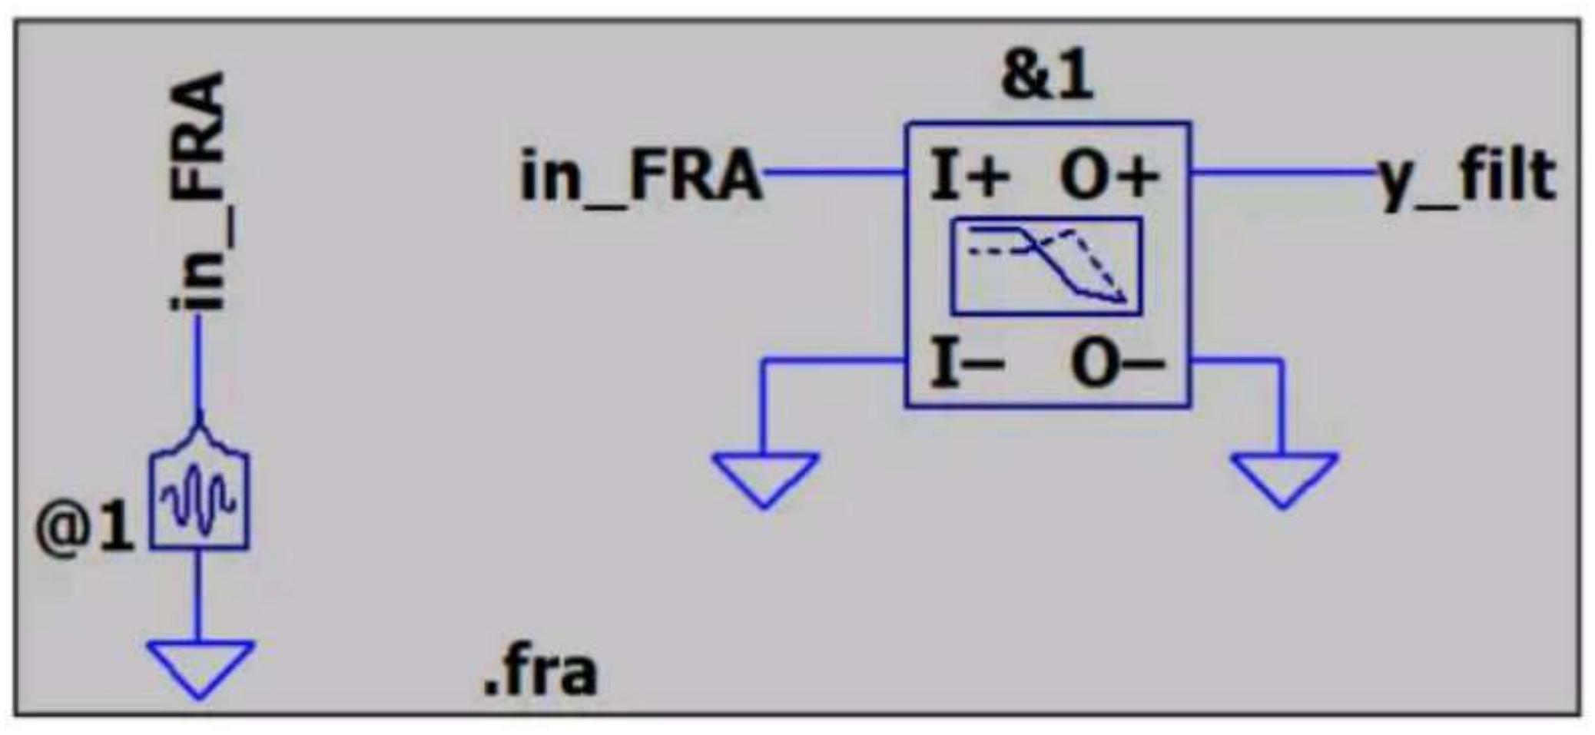

Fig.11: the FRA device and probe as added to the Fig.8 circuit.

Frequency Response Analysis

LTspice provides an alternative to AC analysis called Frequency Response Analysis (FRA).

AC analysis uses a linearised model of the circuit to find the frequency response, which is fine for analog filters and amplifiers, but does not work for switched circuits.

FRA was added to LTspice primarily to assist with the stability analysis of switch-mode power supplies (SMPS). It can help to find the loop gain of a feedback system but it has other uses.

It works by automatically running multiple transient simulations at different frequencies, measuring waveform amplitudes and calculating gains.This means the fully operational circuit is analysed, not just a linearised version, so any switching (or in this case sampling) operations operate normally.

We discussed LTspice's FRA in the March and April 2024 *Circuit Surgery* columns, where we illustrated the basic principles by analysing an op amp circuit feedback loop.

The FRA measurement is controlled by a special FRA device, which is usually inserted into the feedback loop of the circuit. This device then injects signals and measures the response of the loop.

There is no feedback loop in the Fig.8 circuit, so we cannot use the FRA device on its own in this way. We can insert it at the input to inject the test signals, but as it is not in a loop, the response cannot be measured from there.

Fortunately, LTspice provides another special device called an FRA probe that works in conjunction with the FRA device to make the measurements. It is aimed at complex SMPS circuits that the FRA cannot handle on its own. However, it works for a simple input-to-output frequency response analysis.

The FRA probe has two differential inputs – one for the input and one for the output signal, with both used for the response analysis.

To perform the analysis, we add two components: the FRA device and the FRA probe device to the schematic in Fig.8 (remember the whole circuit also includes B1 to B5 from Fig.6).

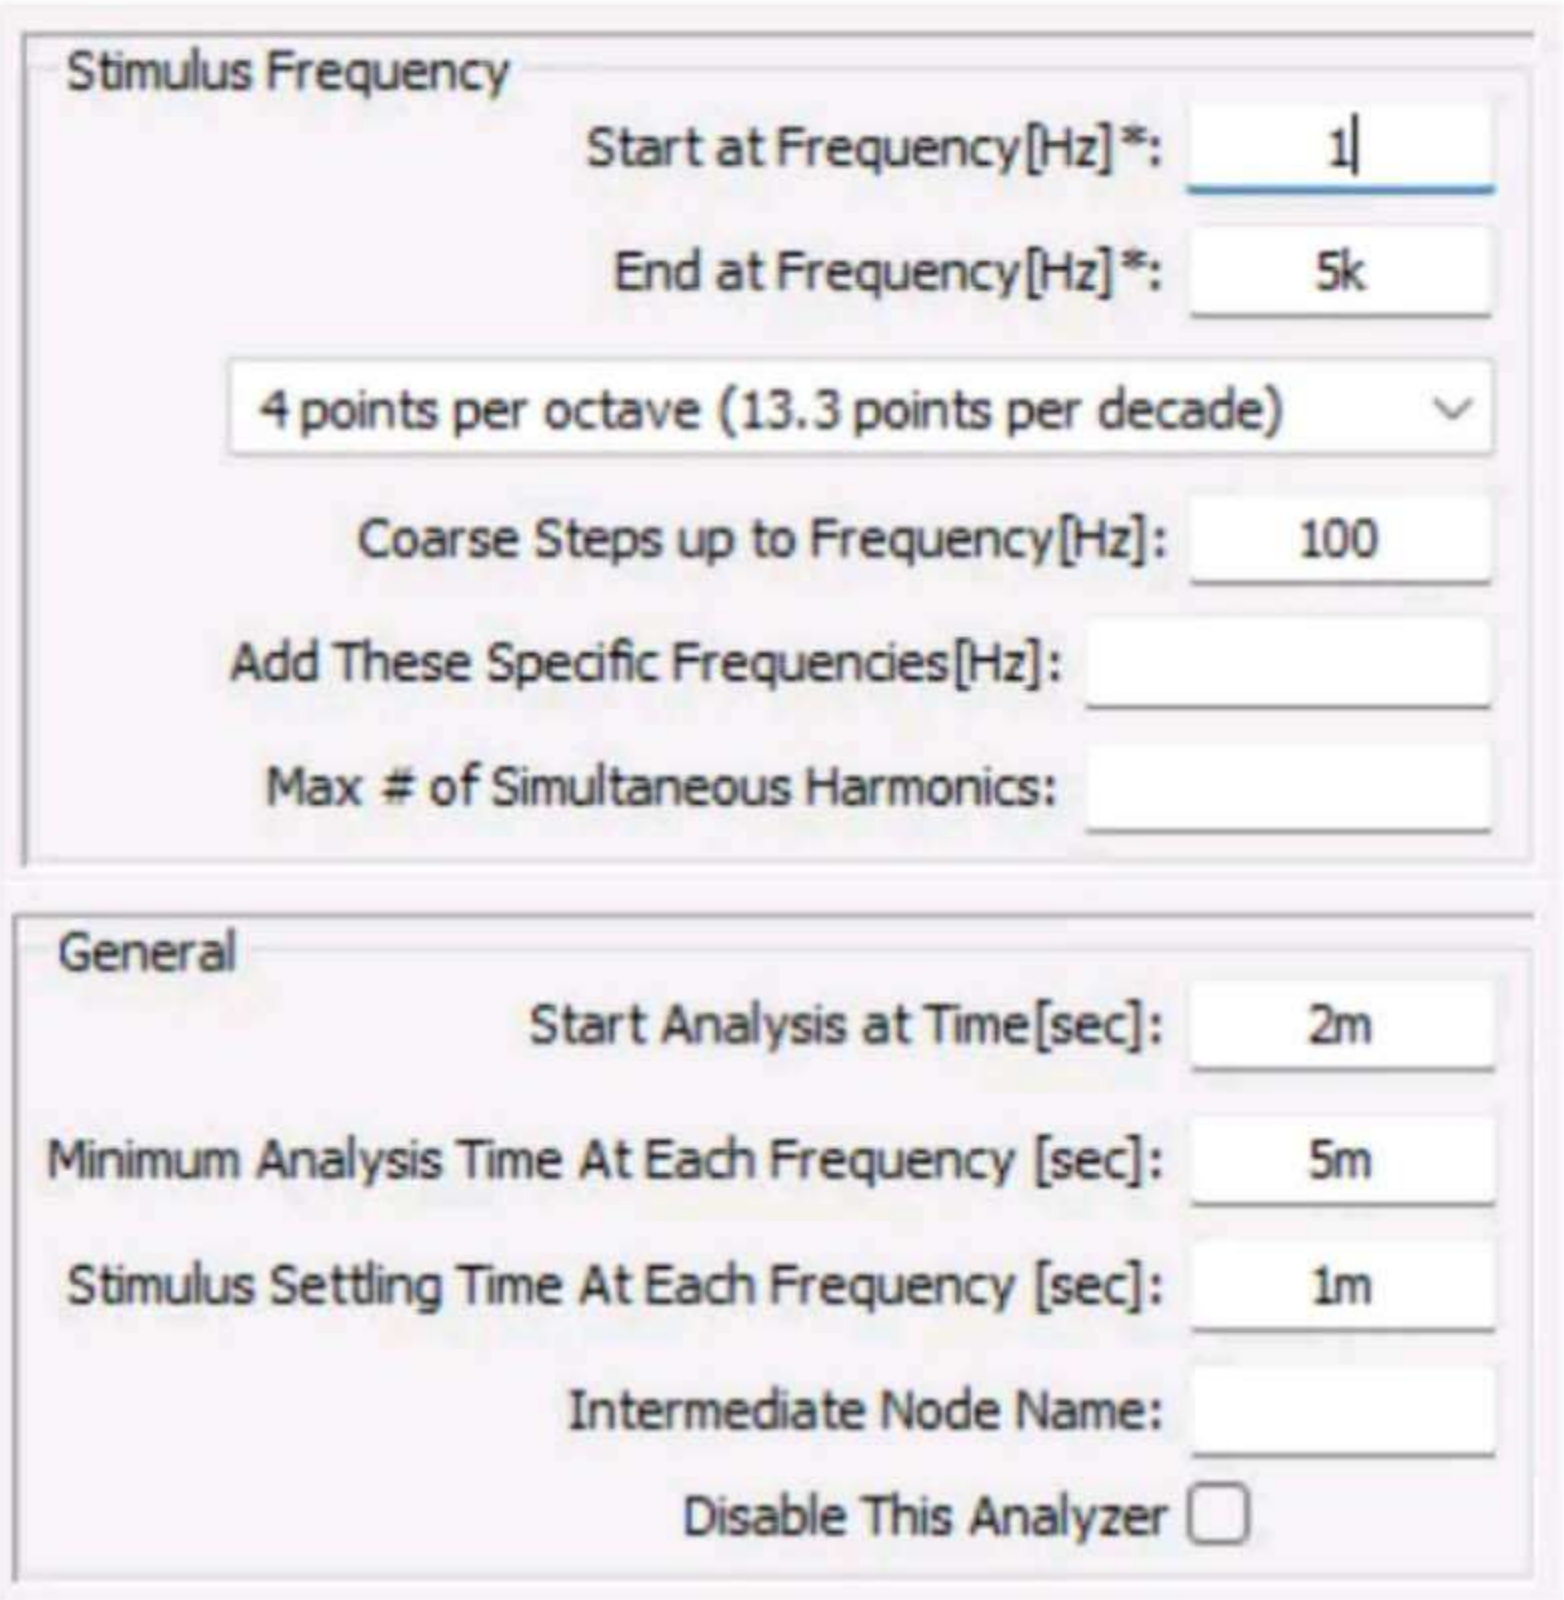

Stimulus Frequency

Start at Frequency[Hz]*: 1

End at Frequency[Hz]*: 5k

4 points per octave (13.3 points per decade)

Coarse Steps up to Frequency[Hz]: 100

Add These Specific Frequencies[Hz]:

Max # of Simultaneous Harmonics:

General

Start Analysis at Time[sec]: 2m

Minimum Analysis Time At Each Frequency [sec]: 5m

Stimulus Settling Time At Each Frequency [sec]: 1m

Intermediate Node Name:

Disable This Analyzer ☐

Fig.12: setting up the FRA stimulus.

This is shown in Fig.11 where the component named @1 is the FRA device injecting the test signals and the component named **&1** is the probe making the measurements.

The input sensing part of the probe is connected to the output of the FRA device on net **in_FRA**. The output sensing part of the probe is connected to the output of the reconstruction filter on net **y_filt**.

The net at the input to the U3 filter is also changed to be **in_FRA** rather than **in** (not shown in Fig.11) to connect the FRA device's test signal to the system input. This leaves the **in** net from V1 not connected to anything, but this is not a problem.

To run the FRA, we need to set the analysis configuration to Transient Frequency response and place the **.fra** directive on the schematic. The way that the FRA operates is configured by right-clicking the FRA device.

Fig.12 shows part of the dialog which appears after doing that. Here, we have set the start and end frequencies for the analysis and selected four points per octave. This controls which frequencies of sinewave will be used in the simulation.

A low start frequency is used to facilitate linear frequency plotting but the "Coarse Steps" setting ensures that not too much time is used simulating low frequency signals. The analysis time settings ensure that enough time is given for the circuit to settle after frequency changes.

If these settings are not used, the results are poor at the high-frequency end

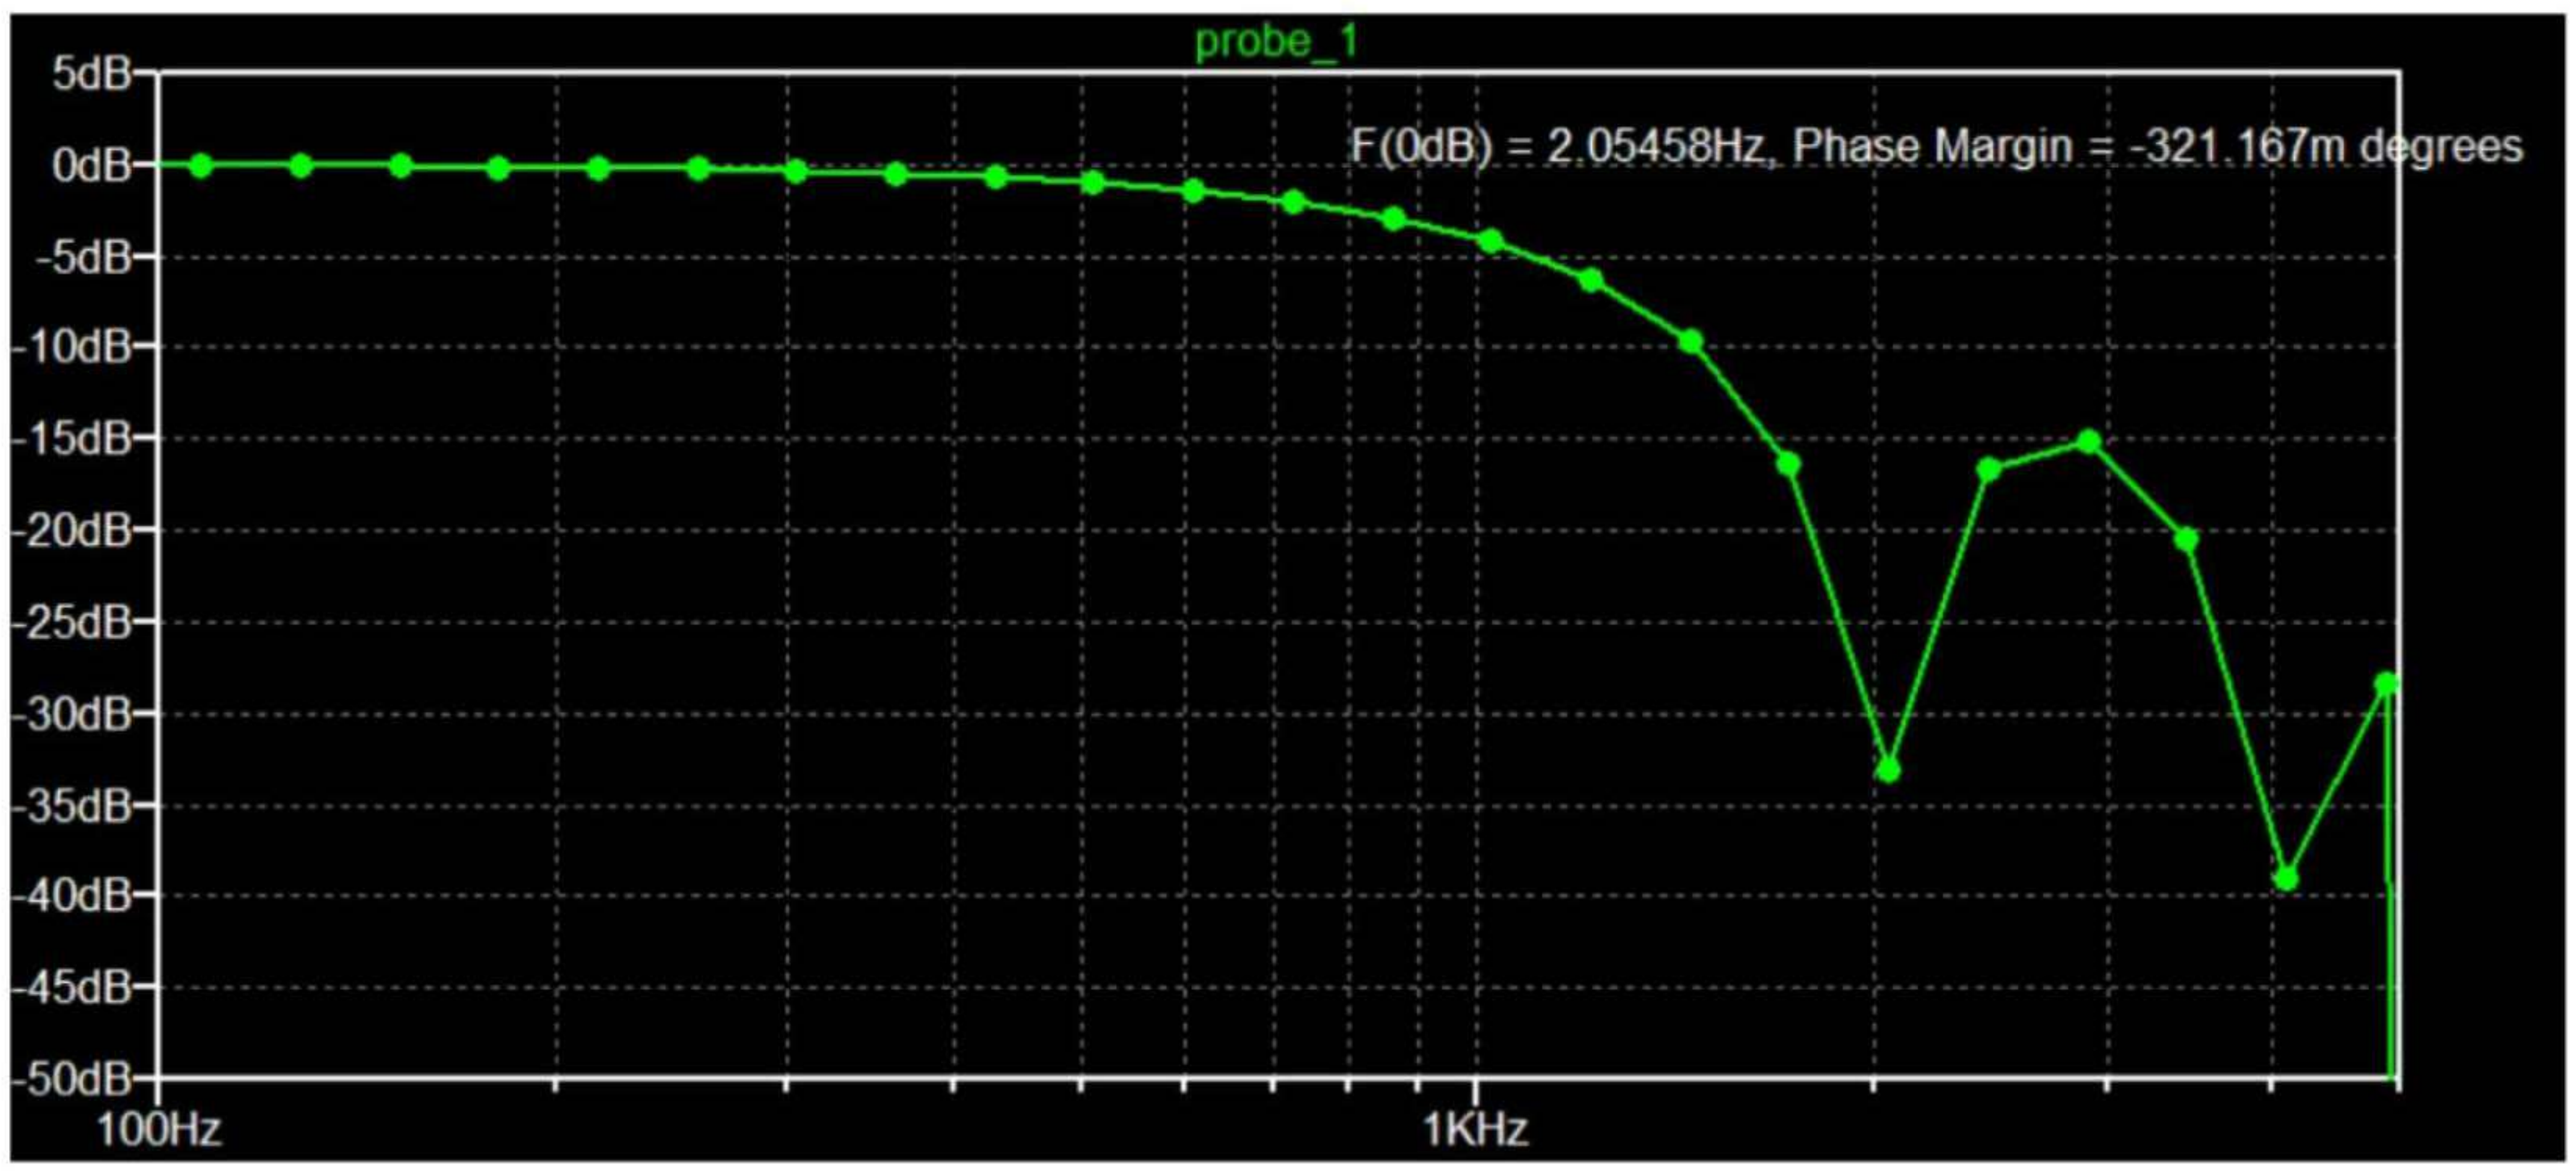

Fig.13: initial results from the FRA.

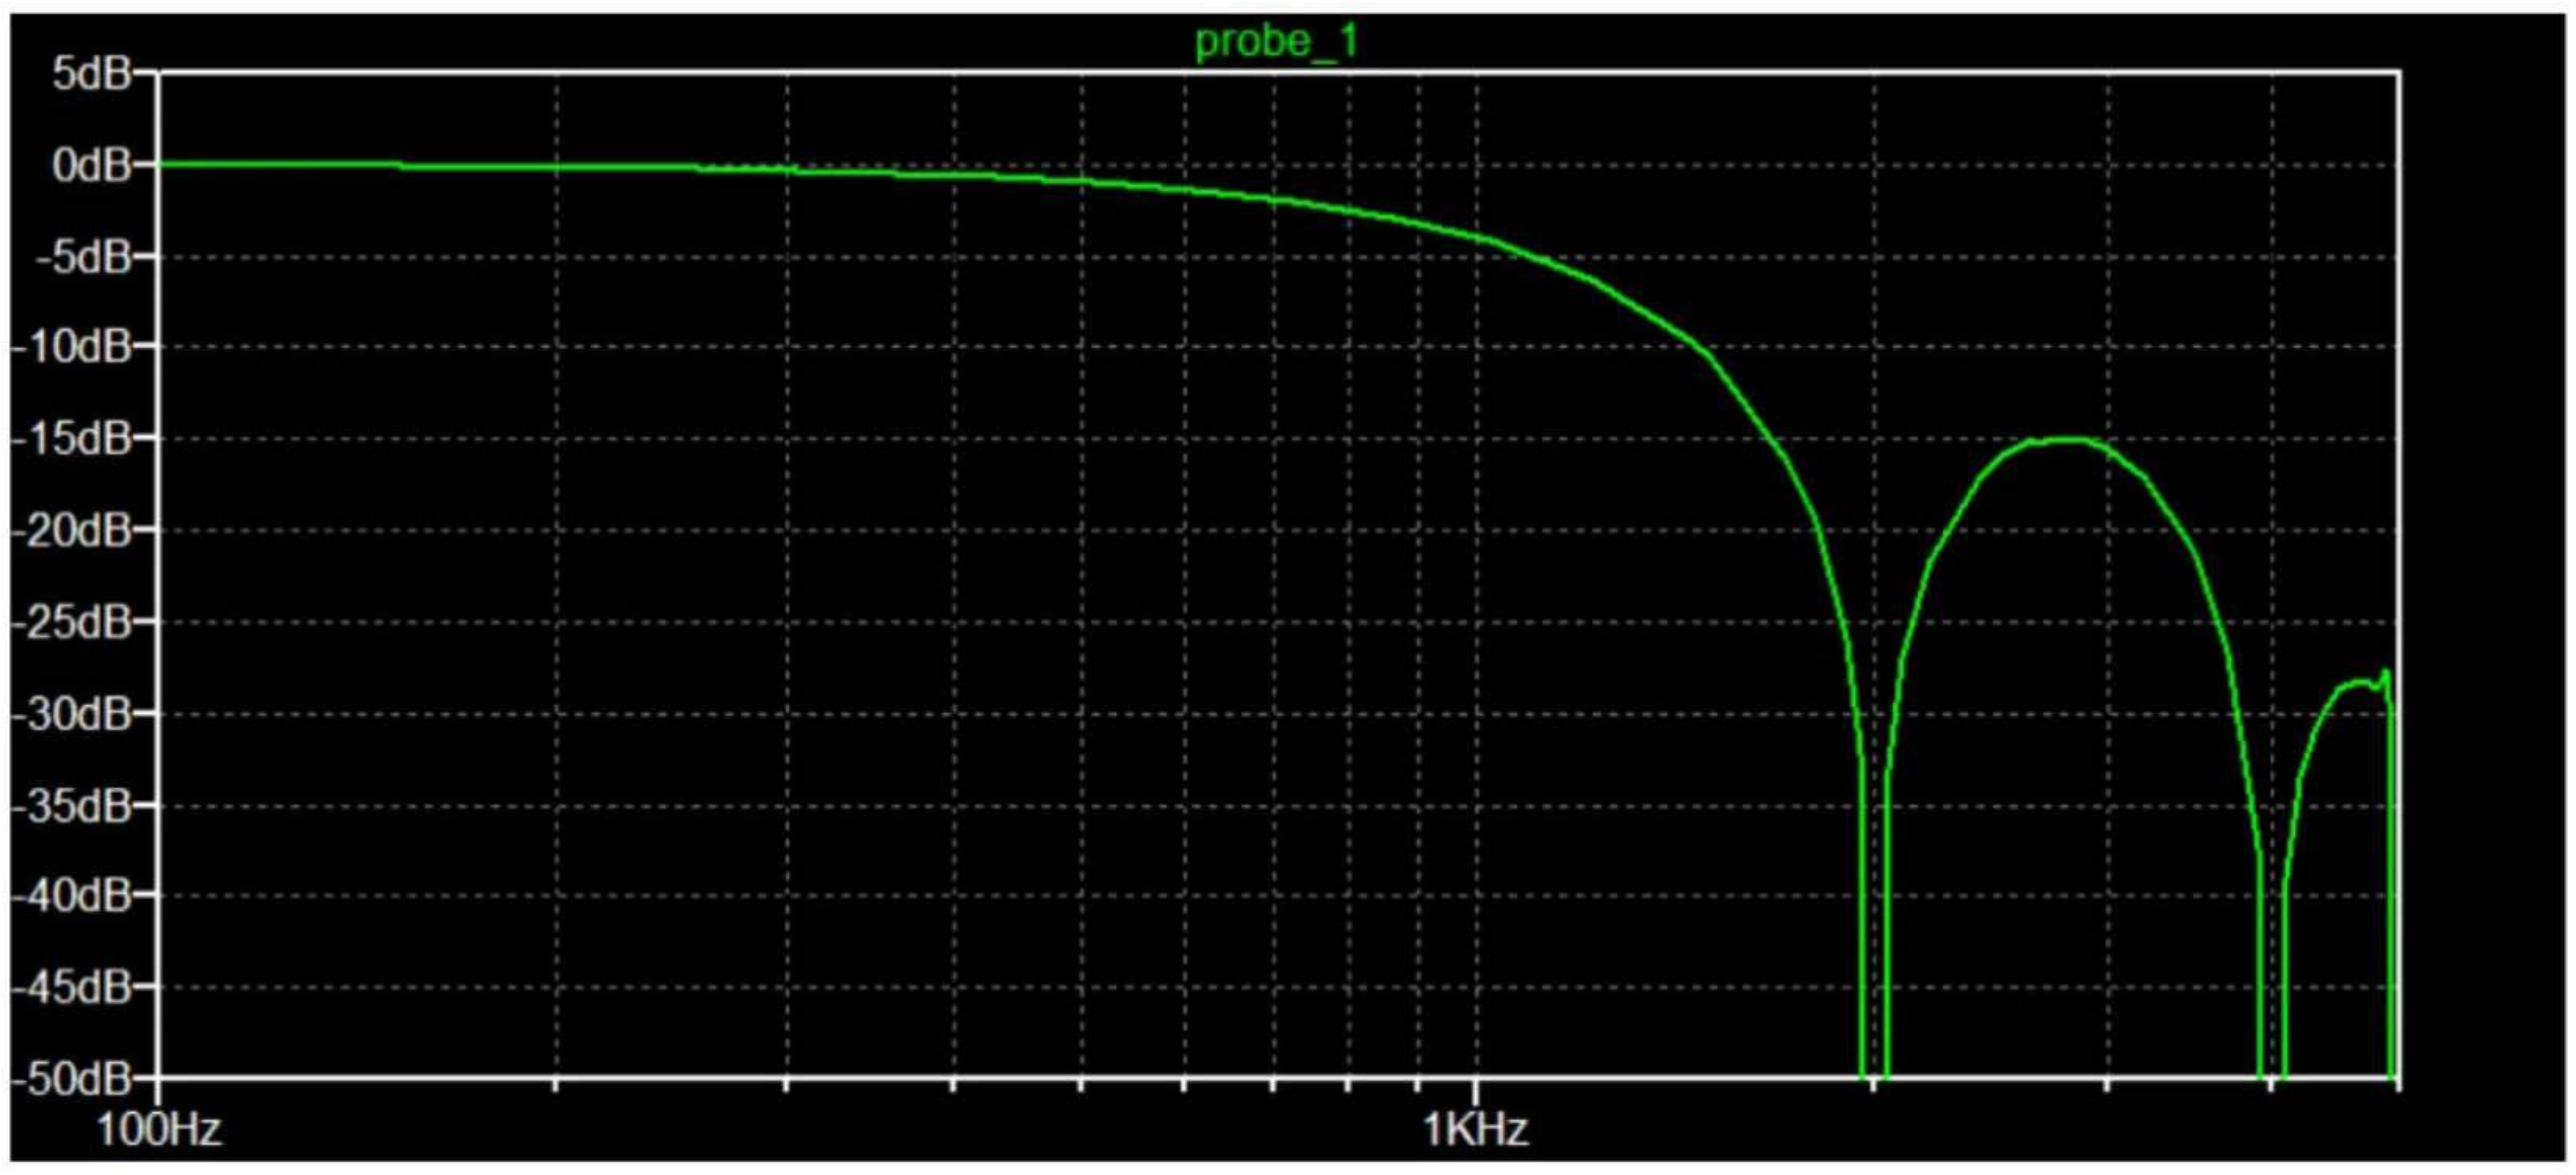

Fig.14: FRA results using more data points.

Simulation files

Most months, LTSpice is used to support descriptions and analysis in Circuit Surgery.
The examples and files are available for download from the PE website:
https://bit.ly/pe-downloads

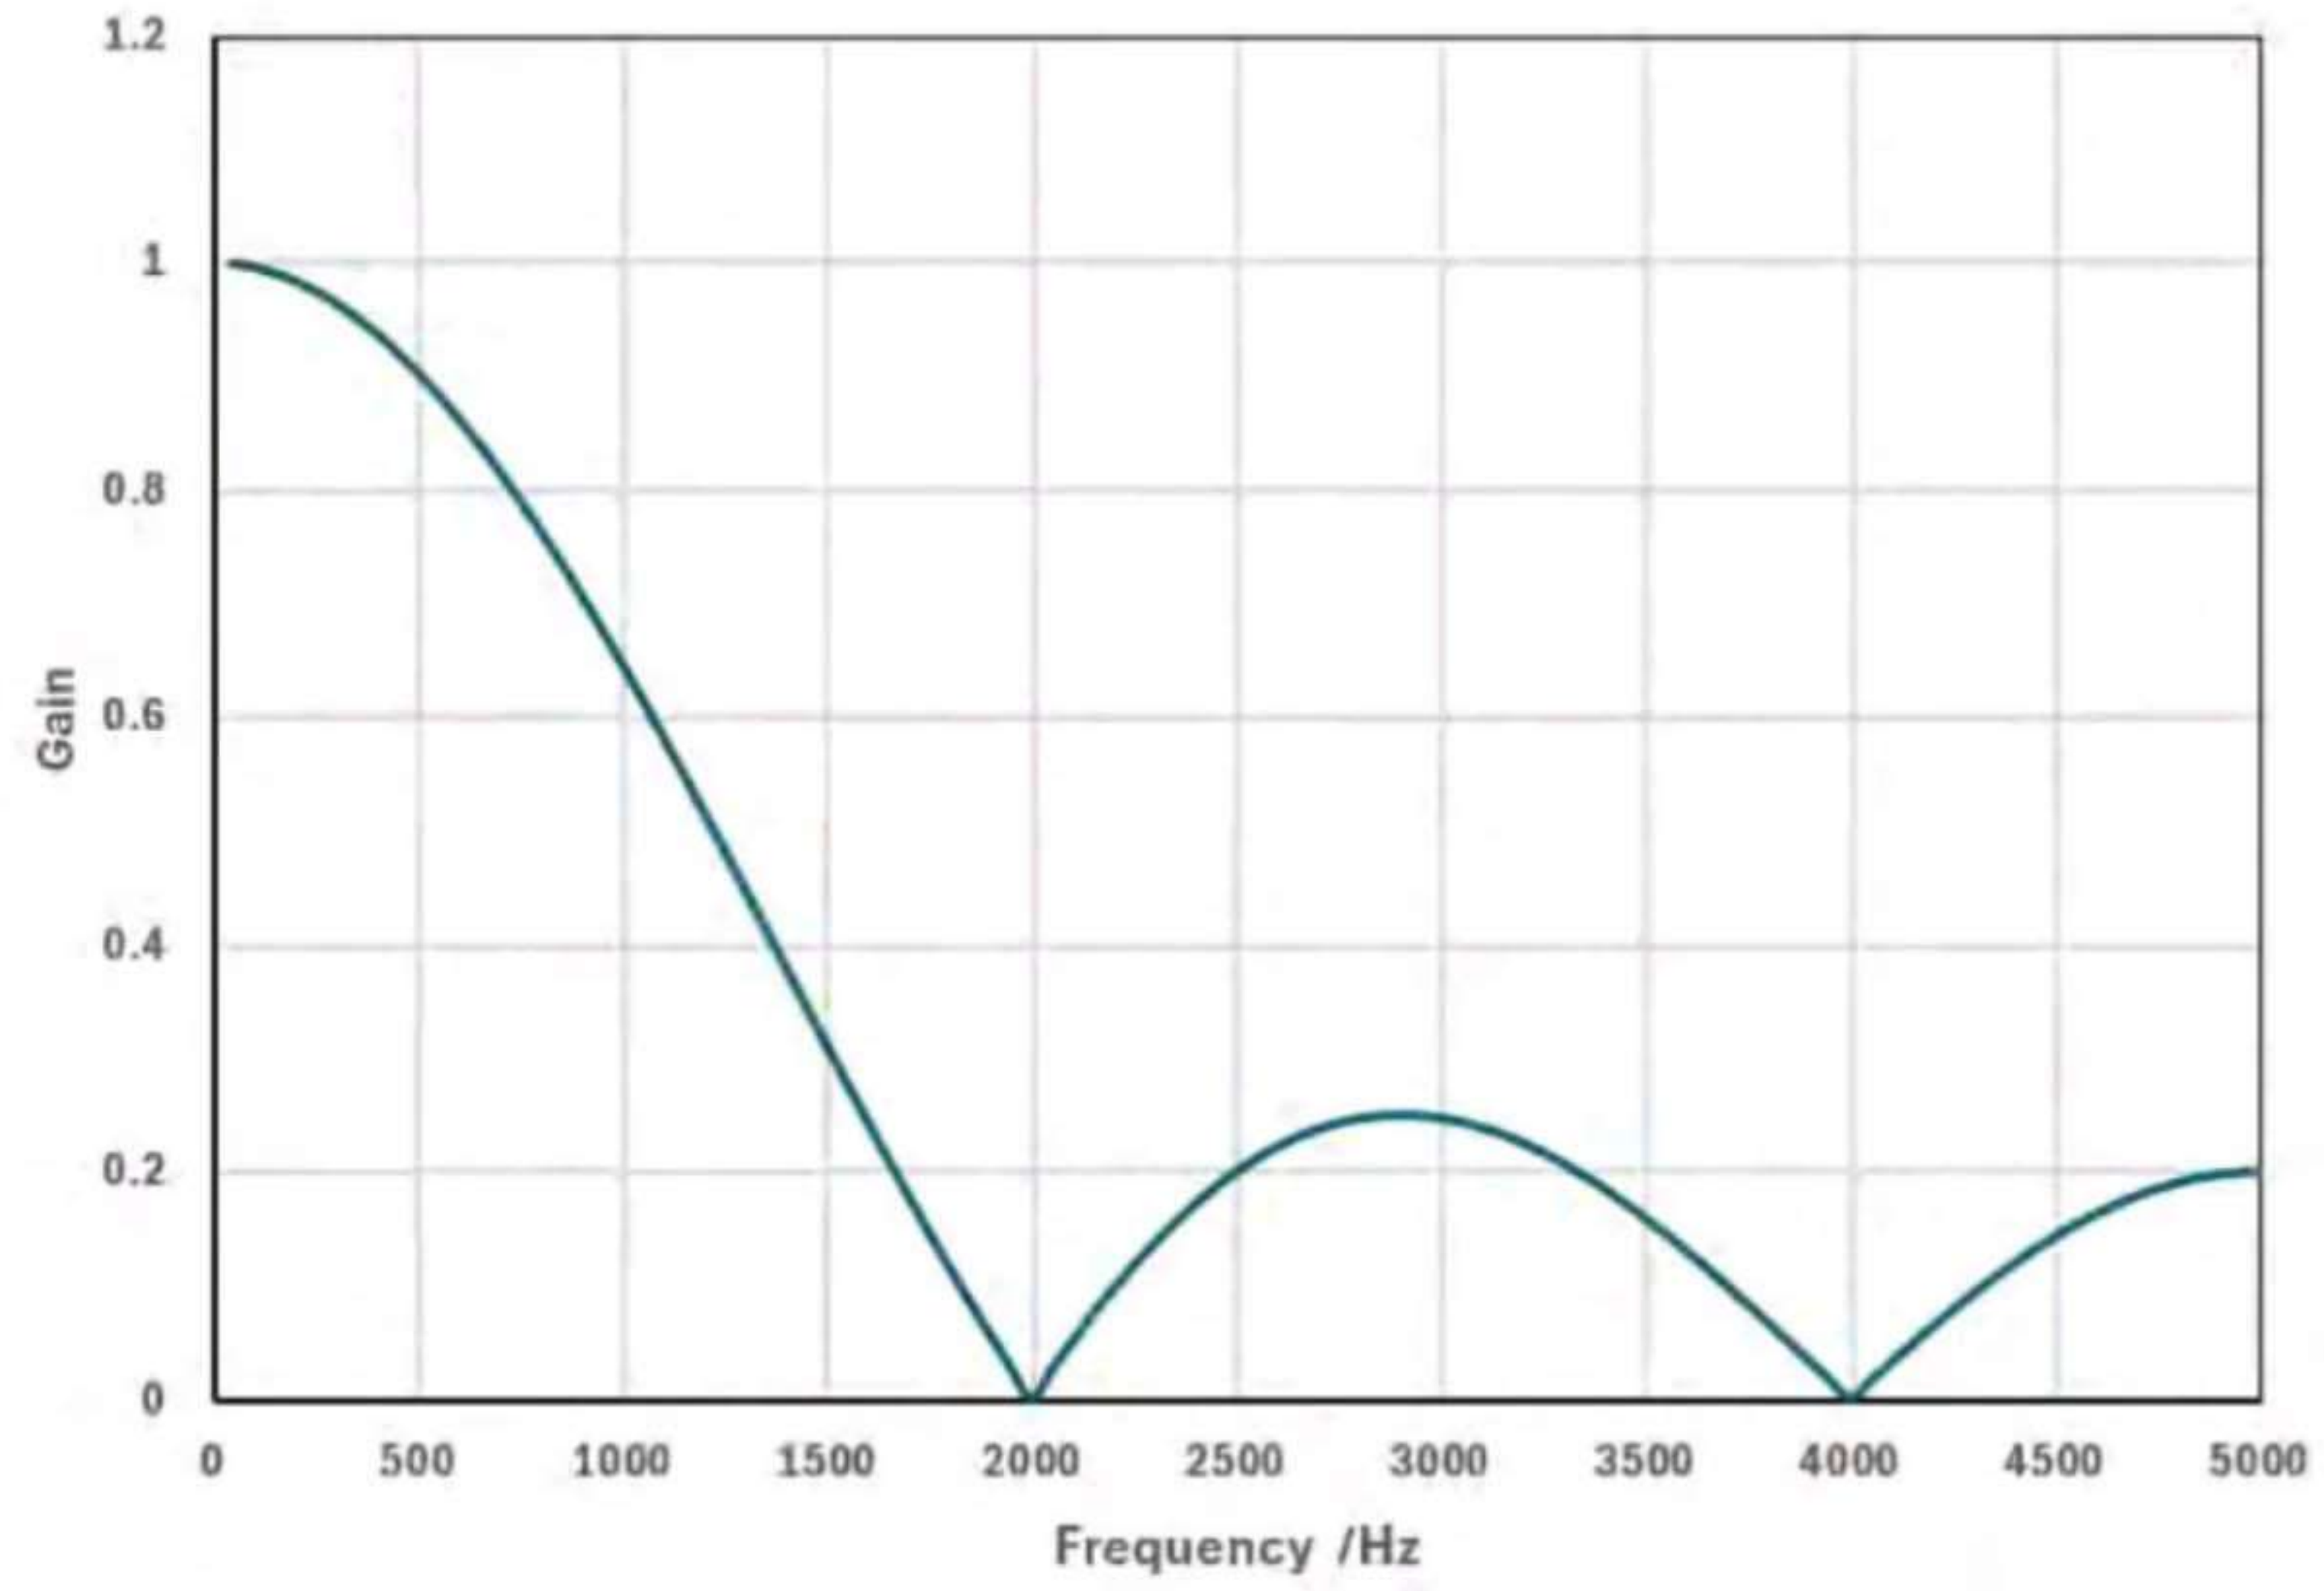

Fig.15: the theoretical response of a five-point moving average filter with a 10kHz sampling frequency.

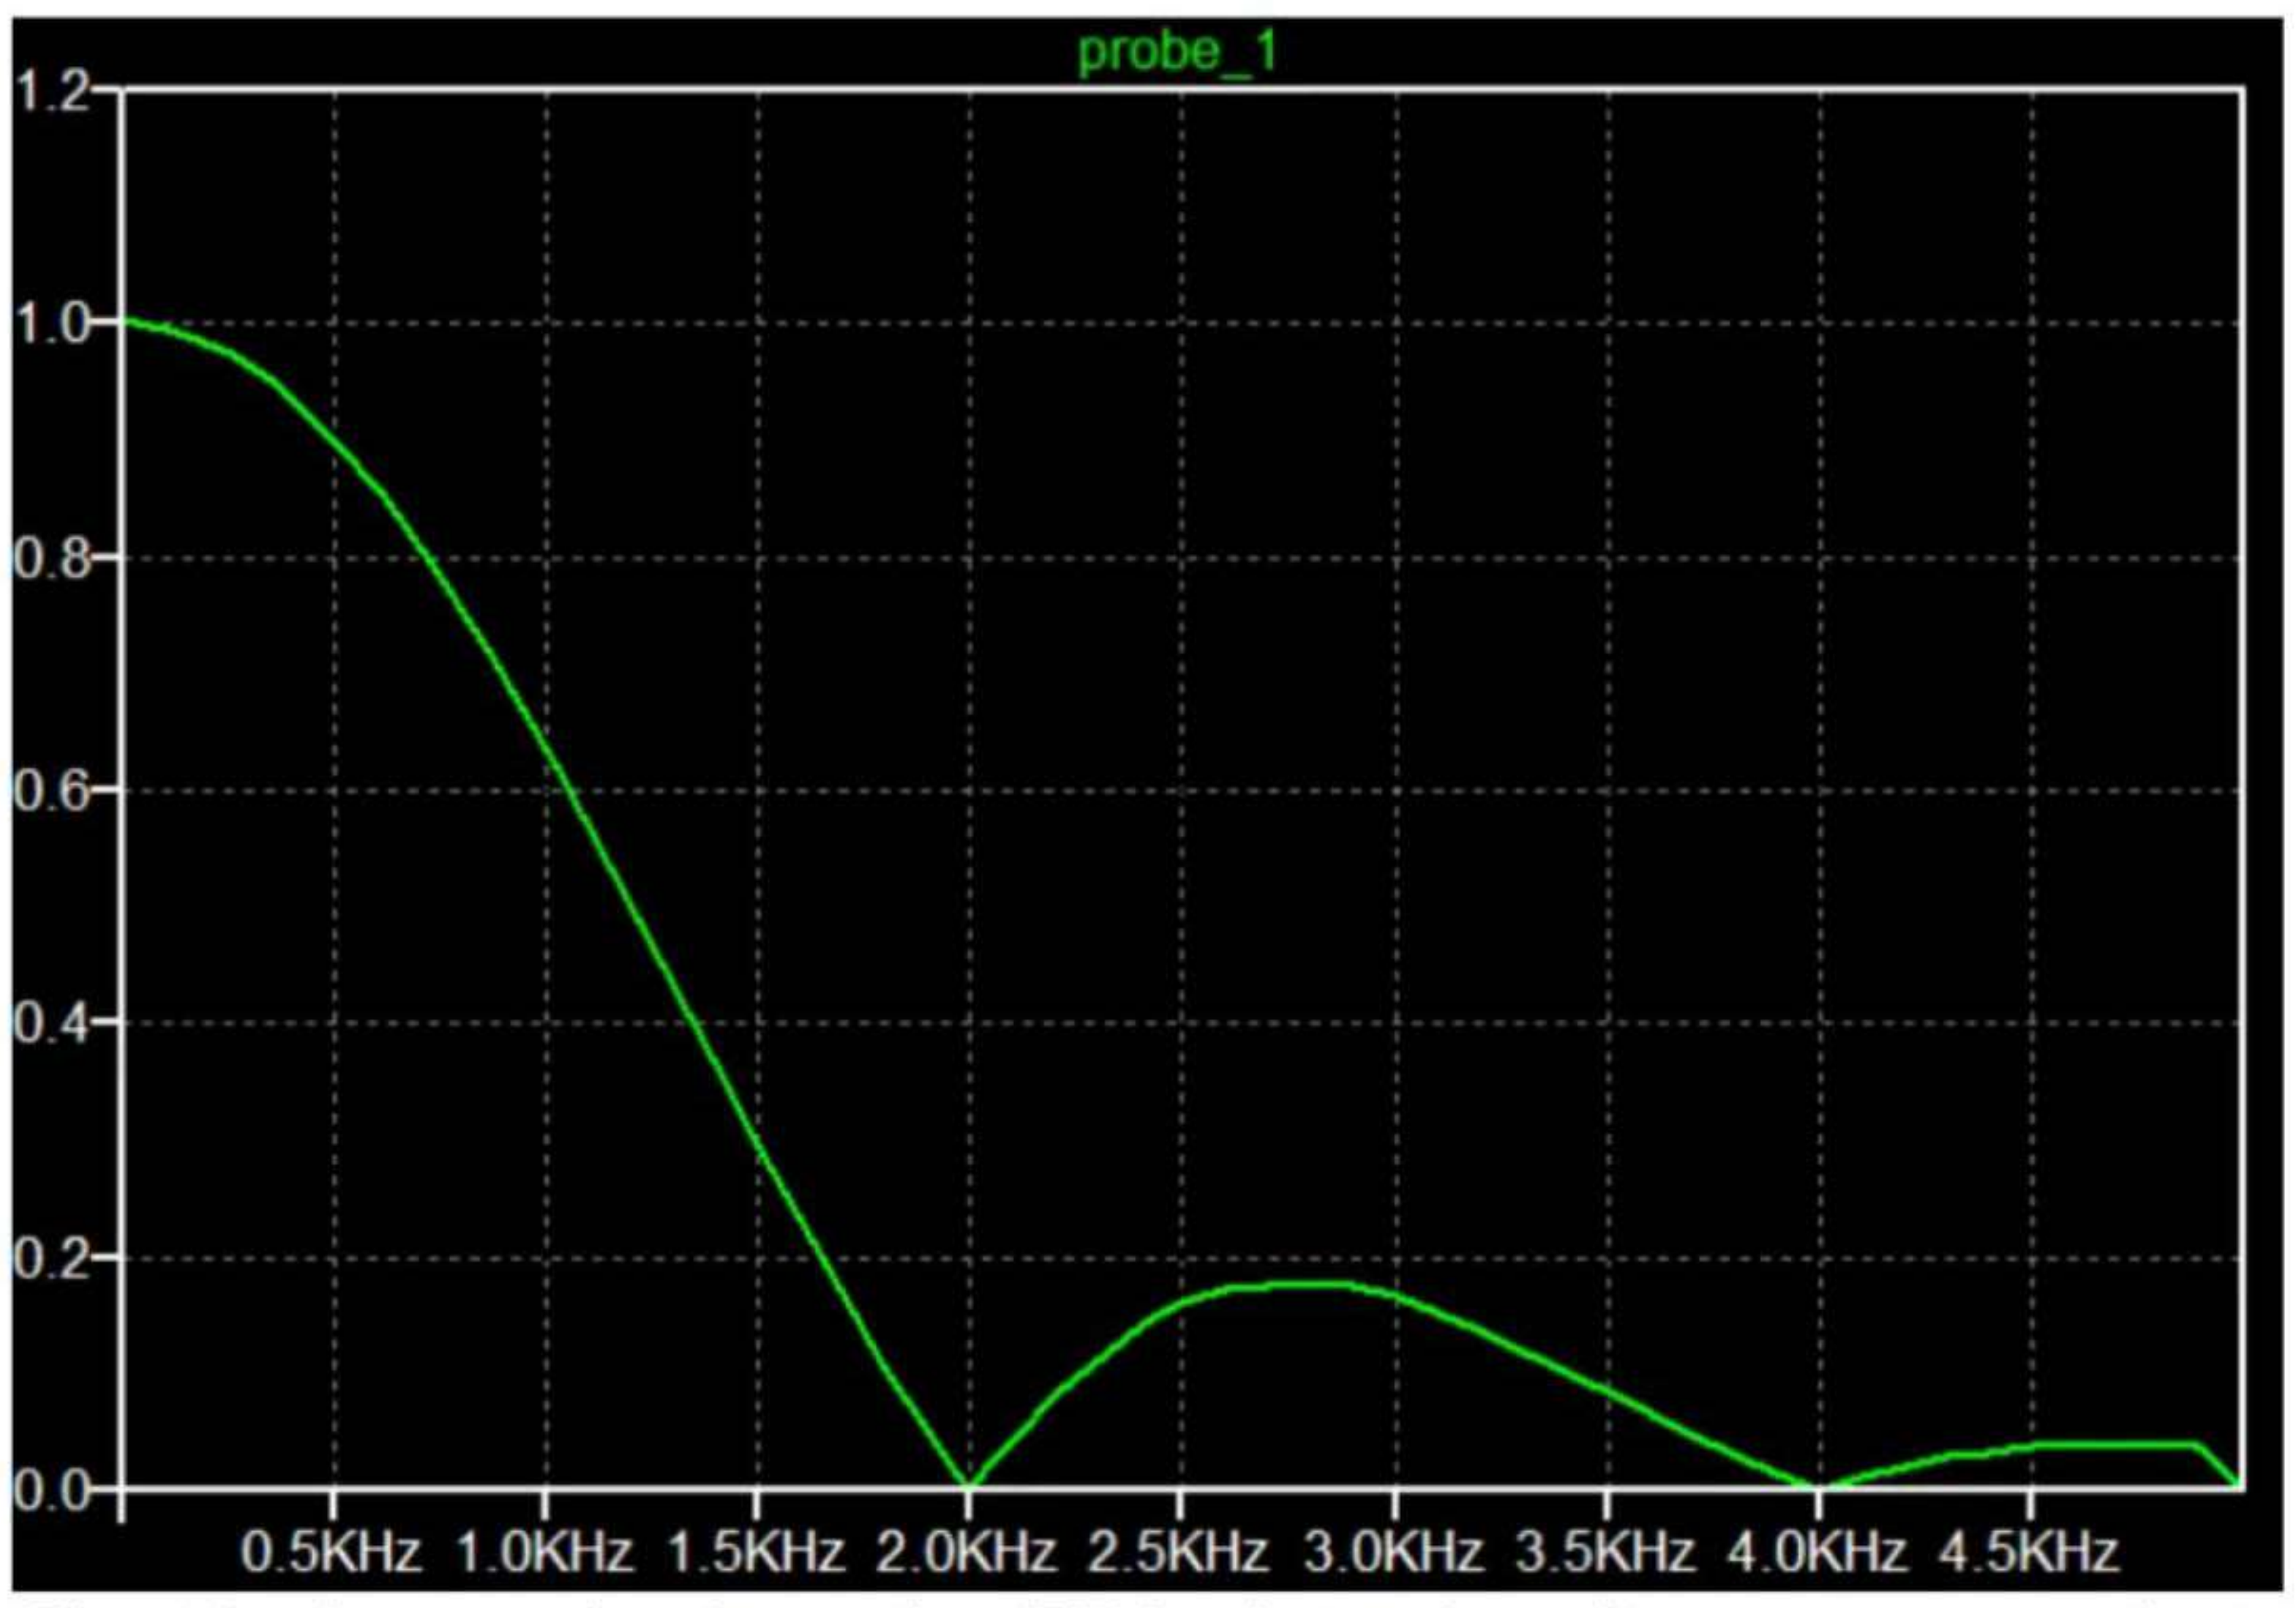

Fig.16: the results from the FRA plotted on linear axes so that we can compare them directly to Fig.15.

of the range, but the values here are not necessarily optimal.

Everything else is left blank (default). For more discussion on the FRA settings, see *Circuit Surgery*, April 2024.

There is a checkbox at the bottom of this dialog to disable the FRA device. That's important as the FRA device must be enabled to run the FRA and disabled when running other types of analysis.

FRA results

The initial results from the FRA (for output `y_filt`) are shown in Fig.13. The text about phase margin is automatically added but is not relevant here – it is easily deleted.

The results are much better than in Fig.10, showing the FRA is working, but the curve lacks detail. This is because the FRA has a maximum of four data points per octave (as per Fig.12), which is insufficient for this graph.

This is a limitation of the FRA in this context, but it can be overcome by adding frequency points in the "Add These Specific Frequencies[Hz]:" box. The improved result in Fig.14 was obtained by adding about 25 extra data points in the 1.5kHz to 4.9kHz range.

The theoretical frequency response (the gain magnitude at frequency f) from a N-sample moving average filter is given by:

$$\left|\frac{\sin\left(N\pi\dfrac{f}{f_s}\right)}{N\sin\left(\pi\dfrac{f}{f_s}\right)}\right|$$

As usual, f_S is the sampling rate. The gain is 1 at $f=0$. This is plotted use Excel in Fig.15, and the FRA results are plotted on linear axes in Fig.16 (the same data as in Fig.14).

The results are similar, but the FRA shows lower gain than the pure moving average function at higher frequencies. That is due to the action of the anti-aliasing and reconstruction filters.

We will continue to look at signal processing next month.

Fluke/Philips PM3092 Oscilloscope
2+2 Channel 200MHz Delay TB,
Autoset etc – £250

STEWART OF READING
17A King Street, Mortimer, near Reading, RG7 3RS
Telephone: **0118 933 1111** Fax: **0118 933 2375**
USED ELECTRONIC TEST EQUIPMENT
Check website **www.stewart-of-reading.co.uk**
(ALL PRICES PLUS CARRIAGE & VAT)
Please check availability before ordering or calling in

HP 54600B Oscilloscope
Analogue/Digital Dual Trace 100MHz
Only £75, with accessories £125

LAMBDA GENESYS	PSU GEN100-15 100V 15A Boxed As New	£400
LAMBDA GENESYS	PSU GEN50-30 50V 30A	£400
IFR 2025	Signal Generator 9kHz – 2.51GHz Opt 04/11	£900
IFR 2948B	Communication Service Monitor Opts 03/25 Avionics	POA
IFR 6843	Microwave Systems Analyser 10MHz – 20GHz	POA
R&S APN62	Syn Function Generator 1Hz – 260kHz	£295
Agilent 8712ET	RF Network Analyser 300kHz – 1300MHz	POA
HP8903A/B	Audio Analyser	£750 – £950
HP8757D	Scaler Network Analyser	POA
HP3325A	Synthesised Function Generator	£195
HP3561A	Dynamic Signal Analyser	£650
HP6032A	PSU 0-60V 0-50A 1000W	£750
HP6622A	PSU 0-20V 4A Twice or 0-50V 2A Twice	£350
HP6624A	PSU 4 Outputs	£400
HP6632B	PSU 0-20V 0-5A	£195
HP6644A	PSU 0-60V 3.5A	£400
HP6654A	PSU 0-60V 0-9A	£500
HP8341A	Synthesised Sweep Generator 10MHz – 20GHz	£2,000
HP83630A	Synthesised Sweeper 10MHz – 26.5 GHz	POA
HP83624A	Synthesised Sweeper 2 – 20GHz	POA
HP8484A	Power Sensor 0.01-18GHz 3nW-10µW	£75
HP8560E	Spectrum Analyser Synthesised 30Hz – 2.9GHz	£1,750
HP8563A	Spectrum Analyser Synthesised 9kHz – 22GHz	£2,250
HP8566B	Spectrum Analsyer 100Hz – 22GHz	£1,200
HP8662A	RF Generator 10kHz – 1280MHz	£750
Marconi 2022E	Synthesised AM/FM Signal Generator 10kHz – 1.01GHz	£325
Marconi 2024	Synthesised Signal Generator 9kHz – 2.4GHz	£800
Marconi 2030	Synthesised Signal Generator 10kHz – 1.35GHz	£750
Marconi 2023A	Signal Generator 9kHz – 1.2GHz	£700

Marconi 2305	Modulation Meter	£250
Marconi 2440	Counter 20GHz	£295
Marconi 2945/A/B	Communications Test Set Various Options	POA
Marconi 2955	Radio Communications Test Set	£595
Marconi 2955A	Radio Communications Test Set	£725
Marconi 2955B	Radio Communications Test Set	£800
Marconi 6200	Microwave Test Set	£1,500
Marconi 6200A	Microwave Test Set 10MHz – 20GHz	£1,950
Marconi 6200B	Microwave Test Set	£2,300
Marconi 6960B	Power Meter with 6910 sensor	£295
Tektronix TDS3052B	Oscilloscope 500MHz 2.5GS/s	£1,250
Tektronix TDS3032	Oscilloscope 300MHz 2.5GS/s	£995
Tektronix TDS3012	Oscilloscope 2 Channel 100MHz 1.25GS/s	£450
Tektronix 2430A	Oscilloscope Dual Trace 150MHz 100MS/s	£350
Tektronix 2465B	Oscilloscope 4 Channel 400MHz	£600
Farnell AP60/50	PSU 0-60V 0-50A 1kW Switch Mode	£300
Farnell XA35/2T	PSU 0-35V 0-2A Twice Digital	£75
Farnell AP100-90	Power Supply 100V 90A	£900
Farnell LF1	Sine/Sq Oscillator 10Hz – 1MHz	£45
Racal 1991	Counter/Timer 160MHz 9 Digit	£150
Racal 2101	Counter 20GHz LED	£295
Racal 9300	True RMS Millivoltmeter 5Hz – 20MHz etc	£45
Racal 9300B	As 9300	£75
Solartron 7150/PLUS	6½ Digit DMM True RMS IEEE	£65/£75
Solatron 1253	Gain Phase Analyser 1mHz – 20kHz	£600
Solartron SI 1255	HF Frequency Response Analyser	POA
Tasakago TM035-2	PSU 0-35V 0-2A 2 Meters	£30
Thurlby PL320QMD	PSU 0-30V 0-2A Twice	£160 – £200
Thurlby TG210	Function Generator 0.002-2MHz TTL etc Kenwood Badged	£65

HP/Agilent HP 34401A Digital
Multimeter 6½ Digit £325 – £375

HP33120A	Function Generator 100 microHz – 15MHz	£350
HP53131A	Universal Counter 3GHz Boxed unused	£600
HP53131A	Universal Counter 225MHz	£350
Audio Precision	SYS2712 Audio Analyser – in original box	POA
Datron 4708	Autocal Multifunction Standard	POA
Druck DPI 515	Pressure Calibrator/Controller	£400
Datron 1081	Autocal Standards Multimeter	POA
ENI 325LA	RF Power Amplifier 250kHz – 150MHz 25W 50dB	POA
Keithley 228	Voltage/Current Source	POA
Time 9818	DC Current & Voltage Calibrator	POA

Marconi 2955B Radio
Communications Test Set – £800

PART 1 OF JOHN CLARKE'S

SECURE REMOTE SWITCH

SECURE REMOTE SWITCH RECEIVER
Power
Relay
(Ack.)
Local
TIMER
ON
OFF
SECURE REMOTE SWITCH TRANSMITTER

This UHF remote switch uses a secure rolling code system. The receiver uses all through-hole parts, fits in a compact plastic case and can be powered from 12V or 24V DC. Up to 16 transmitters can be used per receiver; they fit into handy keyfob cases, and you can use a prebuilt transmitter module or discrete components.

This project comprises a receiver unit with a relay and between one and sixteen miniature keyfob transmitters, each having three pushbuttons and an internal battery. It uses a rolling code system to allow the transmitters to trigger the relay in the receiver, which is designed with a DC power supply and low-voltage switching in mind. That makes it ideal for applications like a garage door controller.

Other potential applications include gate control, remote operation of door strikes or switching DC-powered appliances on and off, such as water pumps, fans, LED lights etc. It is compatible with most 12V or 24V solar power systems or can run from mains power via a suitable supply.

There are two versions of the transmitter: one that uses a prebuilt 433.9MHz transmitter module (operating on the internalional ISM band) and another which is slightly cheaper to build and uses all discrete parts for those who like to 'roll their own'.

The transmitter fits into a nice little keyfob case that we will supply in kits for the transmitters. We'll have kits for the discrete and module-based versions; the discrete kits are complete, while the module-based kits come with everything but the transmitter module, for compliance reasons (you can get that from a local store).

The new transmitters also use small A23 alkaline batteries rather than lithium coin cells; this is mainly due to the design of the cases, but it has the advantage that the quality of A23 alkaline batteries is more consistent than lithium coin cells. This also avoids the serious ingestion hazard that coin cells pose for small children.

The discrete transmitter circuit is based on the Remote Control Range Extender from the January 2023 issue. However, in that design, tiny components were used (some as small as 0.6 × 0.3mm!), which made it a real challenge to assemble, even for us. This time, we have used much larger components that are easier to solder,

Transmitter

- Professional keyfob enclosure
- Secure rolling code communication
- Up to 16 transmitters per receiver
- Powered by a 12V 55mAh A23 battery, giving more than two years of life with typical use
- Range: 22m line-of-sight
- Standby current: typically 3μA (26mAh/year)
- Transmitting current: 10mA average over 1s (2.77μAh per transmission)
- Registration current: 10mA average over 2.75s (7.6μAh per registration)
- Transmission rate: 976.5 bits/s (1.024ms per bit)
- Data encoding: Manchester code with a transmission time of 82ms
- Unique code generation: secure UHF rolling code control with 48-bit seed, 24-bit multiplier and 8-bit increment value

Receiver

- 12V or 24V DC operation
- Supply current: 15mA with relay off, 45mA with relay on
- Relay contact rating: 10A (can handle up to 60V DC/42V AC)
- Relay-on timer range: 250ms to 4.5h (see Tables 1 & 4)

Secure Remote Switch, part one

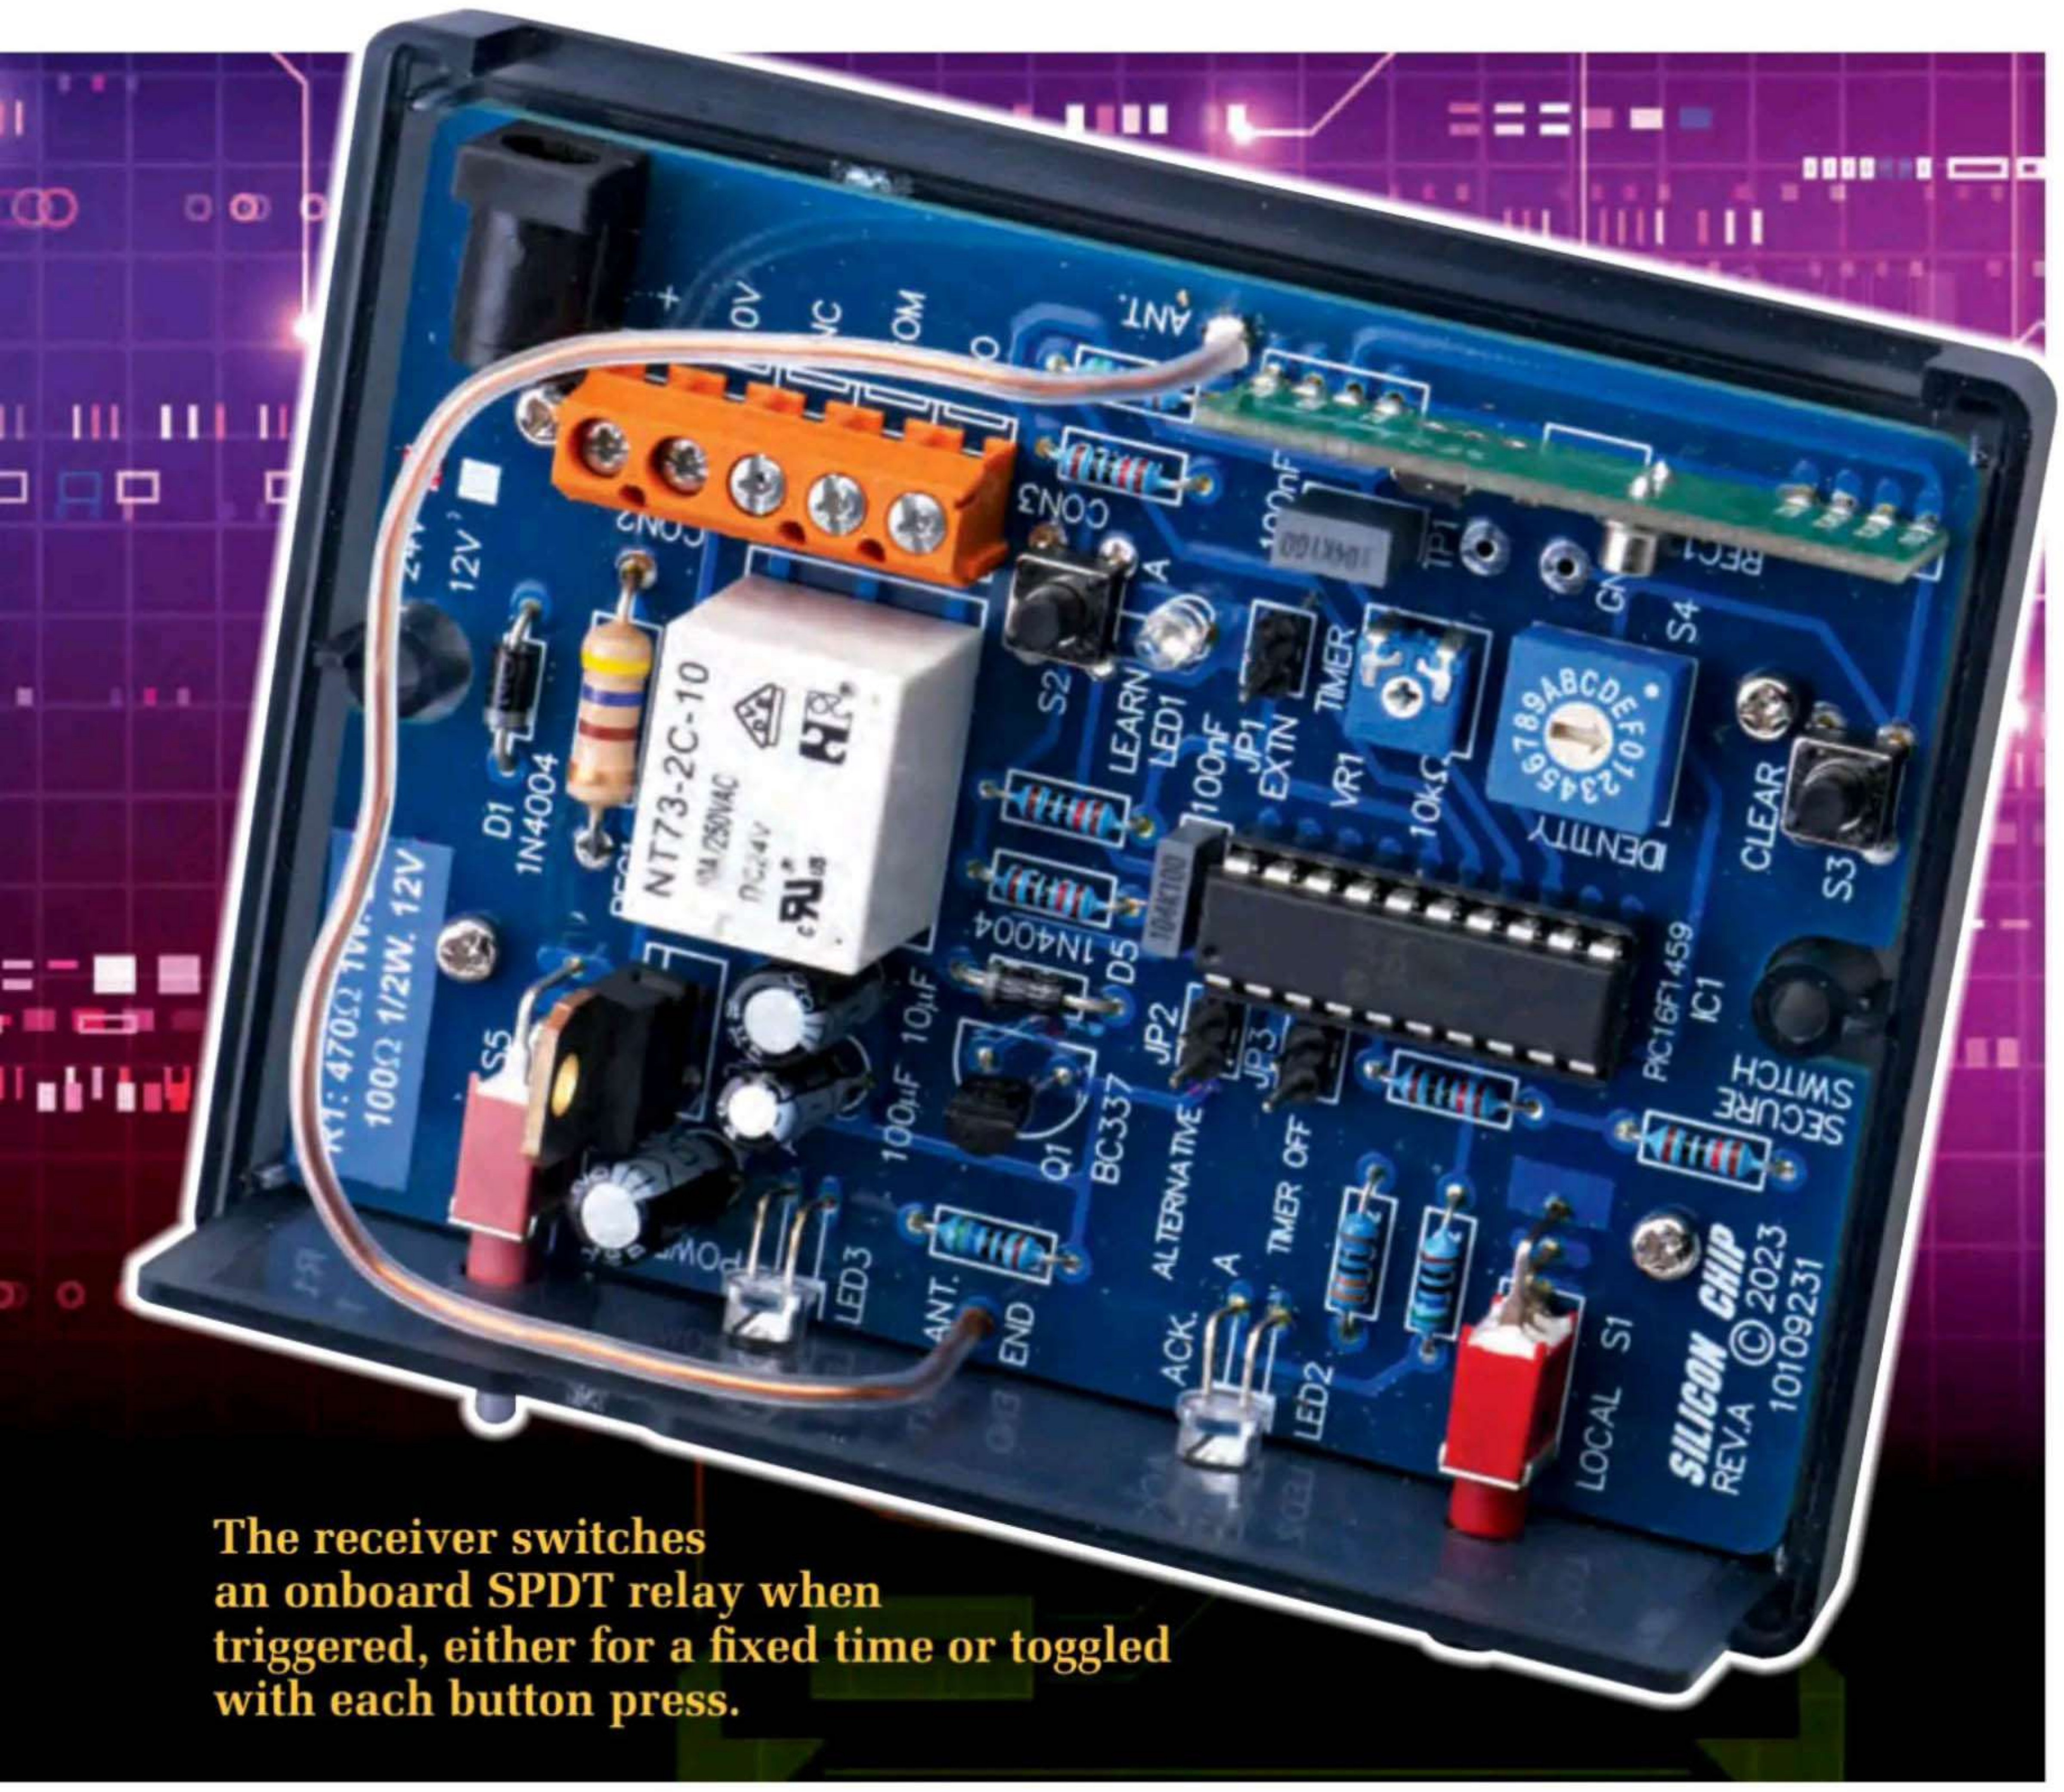

The receiver switches an onboard SPDT relay when triggered, either for a fixed time or toggled with each button press.

so only modest soldering skills are required.

Low-voltage switching

This design can only directly switch low voltages. If you require a remote switch that controls mains voltage, you can use the onboard relay in this design to switch 12V or 24V DC to an external mains-rated relay. It will need to be in its own box with suitable mains connectors, wiring and insulation.

We have decided also to offer a short-form kit for the receiver. You'll need to get a handful of parts yourself, like the case and a few switches, but the kit will save you time and effort gathering the parts to build the Secure Remote Switch.

Security

The Secure Remote Switch uses rolling code wireless transmission to ensure security. That makes it very difficult for someone to trigger the relay on the receiver without having one of your registered remotes. So if it is used to trigger remote-controlled doors, gates and door strikes, the security of your home or premises is maintained.

While secure codes are required for security applications, they also ensure that a similar remote control does not inadvertently switch your appliance on or off. This could happen due to someone close by controlling their own equipment. We've experienced spurious operation of security shutters that we think must have been due to someone using a different remote nearby. That's almost impossible with a rolling code system!

Other controls will not operate the Switch because the transmitter and receiver must be paired before they will work together. Additionally, the code sent between the transmitter and receiver changes each time it is used. That thwarts anyone who might try to capture the code and subsequently resend it in an attempt to control the Switch.

Since the captured code immediately becomes obsolete after use, the Switch will not respond if it is repeated.

You can also build more than one Switch without being concerned about interference between them. The unique transmission code ensures that the Switch receiver will not be activated by anything other than one of the paired handheld remote controls.

The remote control code sent by the handheld remote units can be considered an electronic lock similar to a physical key, except that the key and lock combination changes each time it is used. For the Switch, this key is a specific code the transmitter sends to the receiver. It comprises a long digital data sequence sent in a particular order over a set period. The code must be correct for the receiver to respond.

With a fixed remote control code, an intending thief can receive and store the code sent by the remote control and re-transmit it in an attempt to operate the receiver. However, with a rolling code, the reused code will not trigger the receiver because it requires a different code each time.

Each code that's transmitted differs markedly from one transmission to the next. The codes sent are based on an algorithm (calculation) the transmitter and receiver have in common.

An initial seed value is based on a Microchip Unique Identifier (MUI) value in the transmitter IC. This IC produces a unique set of values that is synchronised with the receiver during registration. These values change each time the Switch is used. Since the handheld remote will have a unique identifier different from any other handheld remote, the uniqueness of the code is ensured.

The odds of picking a correct code at random for our rolling code transmitter is one in 2.8 trillion, making any attempt to break the code by sending out guessed codes unrealistic. The code must also be sent at the correct data rate, with the correct start and stop bit codes and other transmission requirements, including data scrambling that changes for each transmission.

Kits available from Silicon Chip

Receiver kit (SC7401, ~£45): comes with the PCB and most onboard components, including a 12V or 24V relay (specify), except the 433.9MHz receiver module RX1. This kit also includes the case.

Discrete transmitter complete kit (SC6836, ~£12): comes with all parts including the case.

Module-based transmitter short-form kit (SC6837, ~£10): comes with all parts except the transmitter module but including the case.

Parts List – Secure Remote Switch (Transmitter)

1 Supertronic PP43 keyfob enclosure
1 A23 12V battery
1 PIC16LF15323-I/SL programmed with 1010923A.HEX, SOIC-14 (IC1)
1 MCP1703-3302E/DB 3.3V low-dropout regulator, SOT-223 (REG1)
  [Farnell 2113888]
1 1N5819 40V 1A schottky diode (D1)
3 SPST two-pin momentary PCB-mount tactile switches (S1-S3)
1 3mm high-brightness red or green LED (LED1)
2 1µF 25V SMD X7R ceramic capacitors, M3216/1206 size
2 100nF 50V SMD X7R ceramic capacitors, M3216/1206 size
1 220Ω 1% SMD resistor, M3216/1206 size
– up to 16 transmitters can be used per receiver

Extra parts for the module-based version
1 double-sided PCB coded 10109232, 29.8 × 39.4mm
1 433.9MHz UHF ASK transmitter module (TX1) **[rfsolutions QAM-TX2-433]**
1 147mm length of 0.8mm enamelled copper wire

Extra parts for the discrete version
1 double-sided PCB coded 10109233, 29.8 × 39.4mm
1 MICRF113YM6 UHF ASK transmitter, SOT-23-6 (IC2) **[Farnell 2810141]**
1 13.56MHz 5 x 3.2mm SMD crystal (X1) **[Farnell 1611805]**
1 470nH SMD inductor, 610MHz SRF, M2012/0805 size (L1)
  [Coilcraft 0805HT-R47TJLB; Farnell 2286517]
1 68nH SMD inductor, 1.7GHz SRF, M1608/0603 size (L2)
  [Coilcraft 0603CS-68NXJLU; Farnell 2286005]
1 1µF 25V SMD X7R ceramic capacitor, M3216/1206 size
2 18pF 50V SMD C0G/NP0 ceramic capacitors, M3216/1206 size
1 12pF 50V SMD C0G/NP0 ceramic capacitor, M3216/1206 size
1 5pF 50V SMD C0G/NP0 ceramic capacitor, M3216/1206 size
1 167mm length of 0.8mm diameter enamelled copper wire

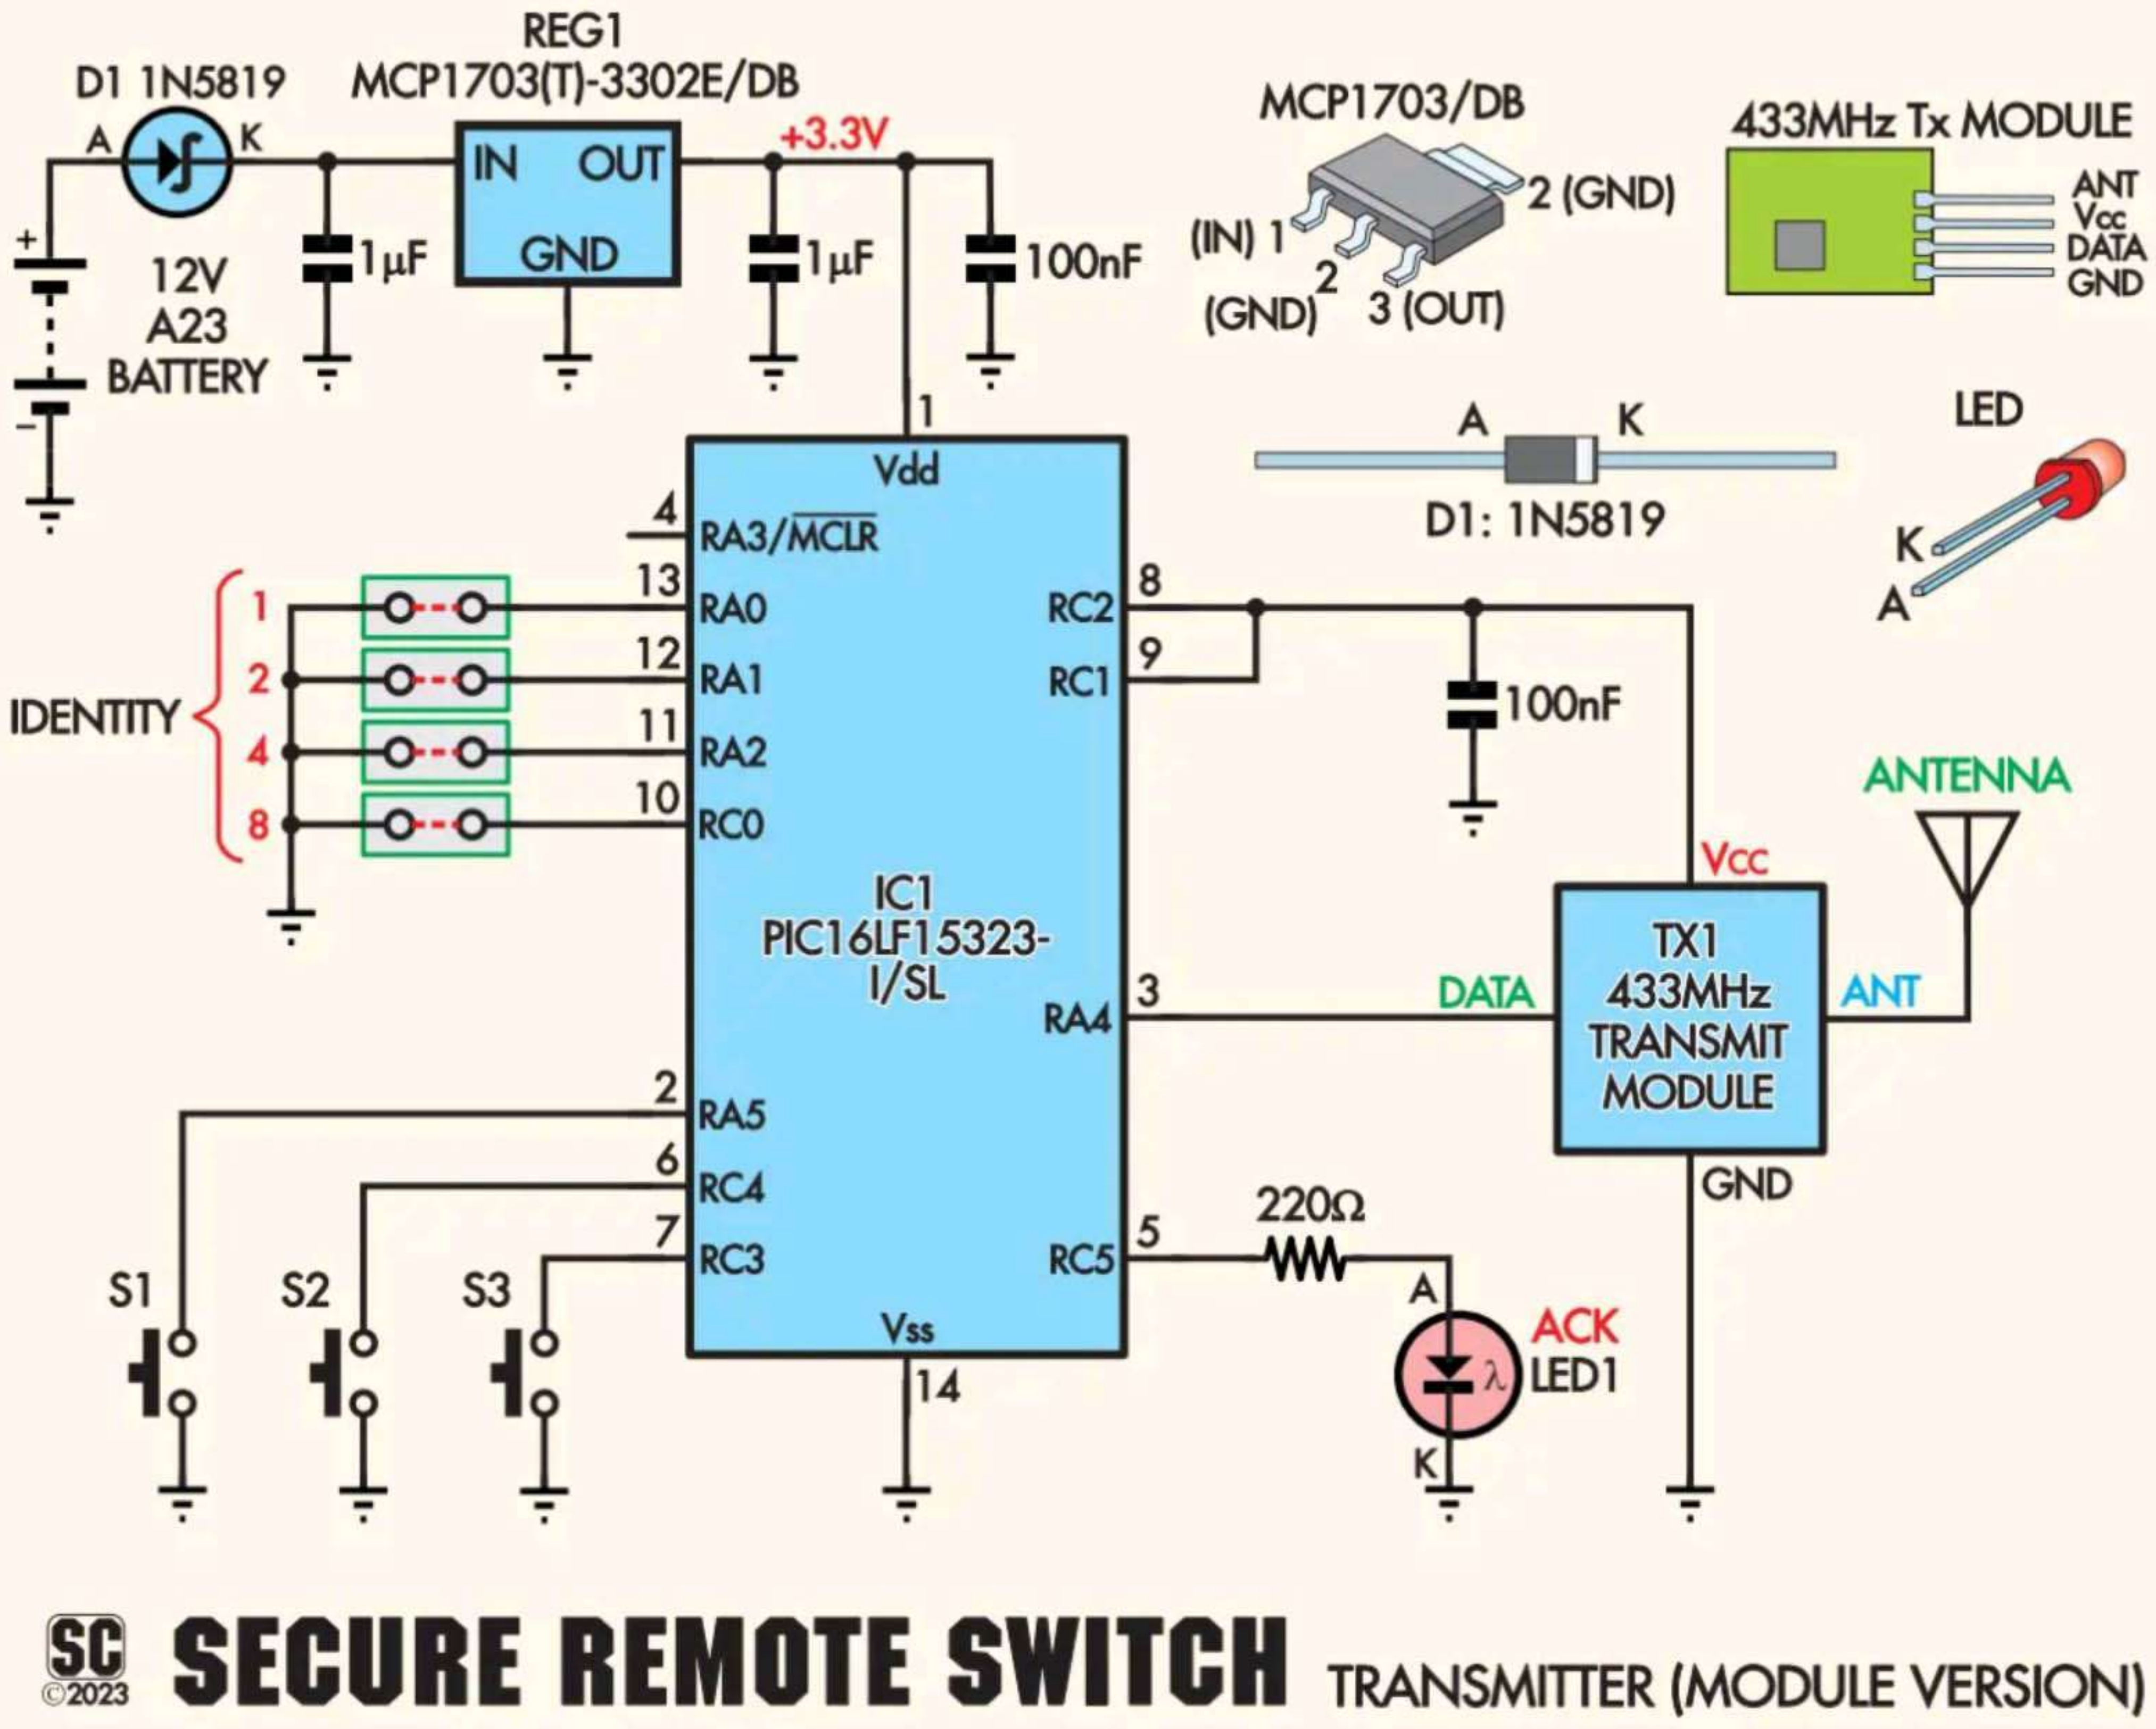

Fig.1: in the module-based transmitter circuit, microcontroller IC1 monitors buttons S1-S3. When one is pressed, it lights LED1, powers up the transmitter module by bringing its pins 8 and 9 high, then produces the ASK data to transmit at its pin 3. When finished, it brings pins 5, 8 and 9 low again and returns to sleep mode.

Other features

Our Switch system has two parts: a professional keyfob-style transmitter and a separate receiver. The keyfob has three pushbutton switches and an acknowledge LED that briefly lights each time one of the switches is pressed. Up to 16 different keyfob transmitters can be used with one receiver.

The receiver has a 10A-rated relay, making it suitable for switching many items. Relays with even higher ratings (eg, 16A) are available if needed. The relay can be controlled by a remote control or a switch on the receiver, and either way, it can be toggled on and off, or switched on for a fixed time. The on-period can be adjusted from 250ms to four-and-a-half hours in two ranges.

Security and registration

Each keyfob transmitter is allocated an Identity number from 0 to 15, set by the positions of coding links on the PCB. Each transmitter is registered to the receiver by sending a synchronising code to the receiver when the receiver is in registration or learning mode.

A facility is included to lock out a particular transmitter after registration. This is useful if a transmitter has been lost. If the lost transmitter is found, it can be easily re-registered. If the identity of the lost transmitter is not known, all transmitters can be locked out, and the ones still in use can be re-registered.

Circuit details

The transmitter circuits are shown in Figs.1 & 2. They have many common parts; each mainly comprises a microcontroller, IC1, and a 433.9MHz UHF transmitter. The UHF transmitter can be either a prebuilt module (Fig.1) or a discrete circuit using a Micrel UHF transmitter IC and associated inductors and capacitors (Fig.2).

Both versions have the same transmission range and fit into the same keyfob enclosure. So which version you wish to build depends on whether you prefer to source the module or solder the discrete parts onto the PCB. The discrete version does have the advantage of potentially being less costly. Both versions utilise a similar wire coil antenna.

The PIC16LF15323 was chosen for IC1 due to its very low standby current and the inclusion of a unique identifier called the Microchip Unique Identifier (MUI). We use the MUI to generate a unique rolling code sequence for each IC; no two transmitters will have the same sequence.

IC1 is usually kept in sleep mode with its internal oscillator stopped and most of its internal circuitry switched off.

Switches S1, S2 and S3 connect to the RA5, RC4 and RC3 digital inputs of IC1, which have internal pullup currents enabled, so those pins are usually high but are pulled low when a button is pressed.

The Identity links (1, 2, 4 & 8) connect to the RA0, RA1, RA2 and RC0 digital inputs, respectively. These are used to differentiate between multiple transmitters used with a given receiver. If only one transmitter is used, it can be set to Identity 0, so none of the Identity pins need to be connected to ground.

At power-up, each Identity input is held high by pullup currents/resistors (within IC1) to the 3.3V rail, similar to the pushbutton inputs. The software then switches off the pullup current for any identity input that is found to be at a low level. That prevents the IC from continuously sourcing current from those pins, which would otherwise add some 25-200µA battery draw per Identity input that's tied low.

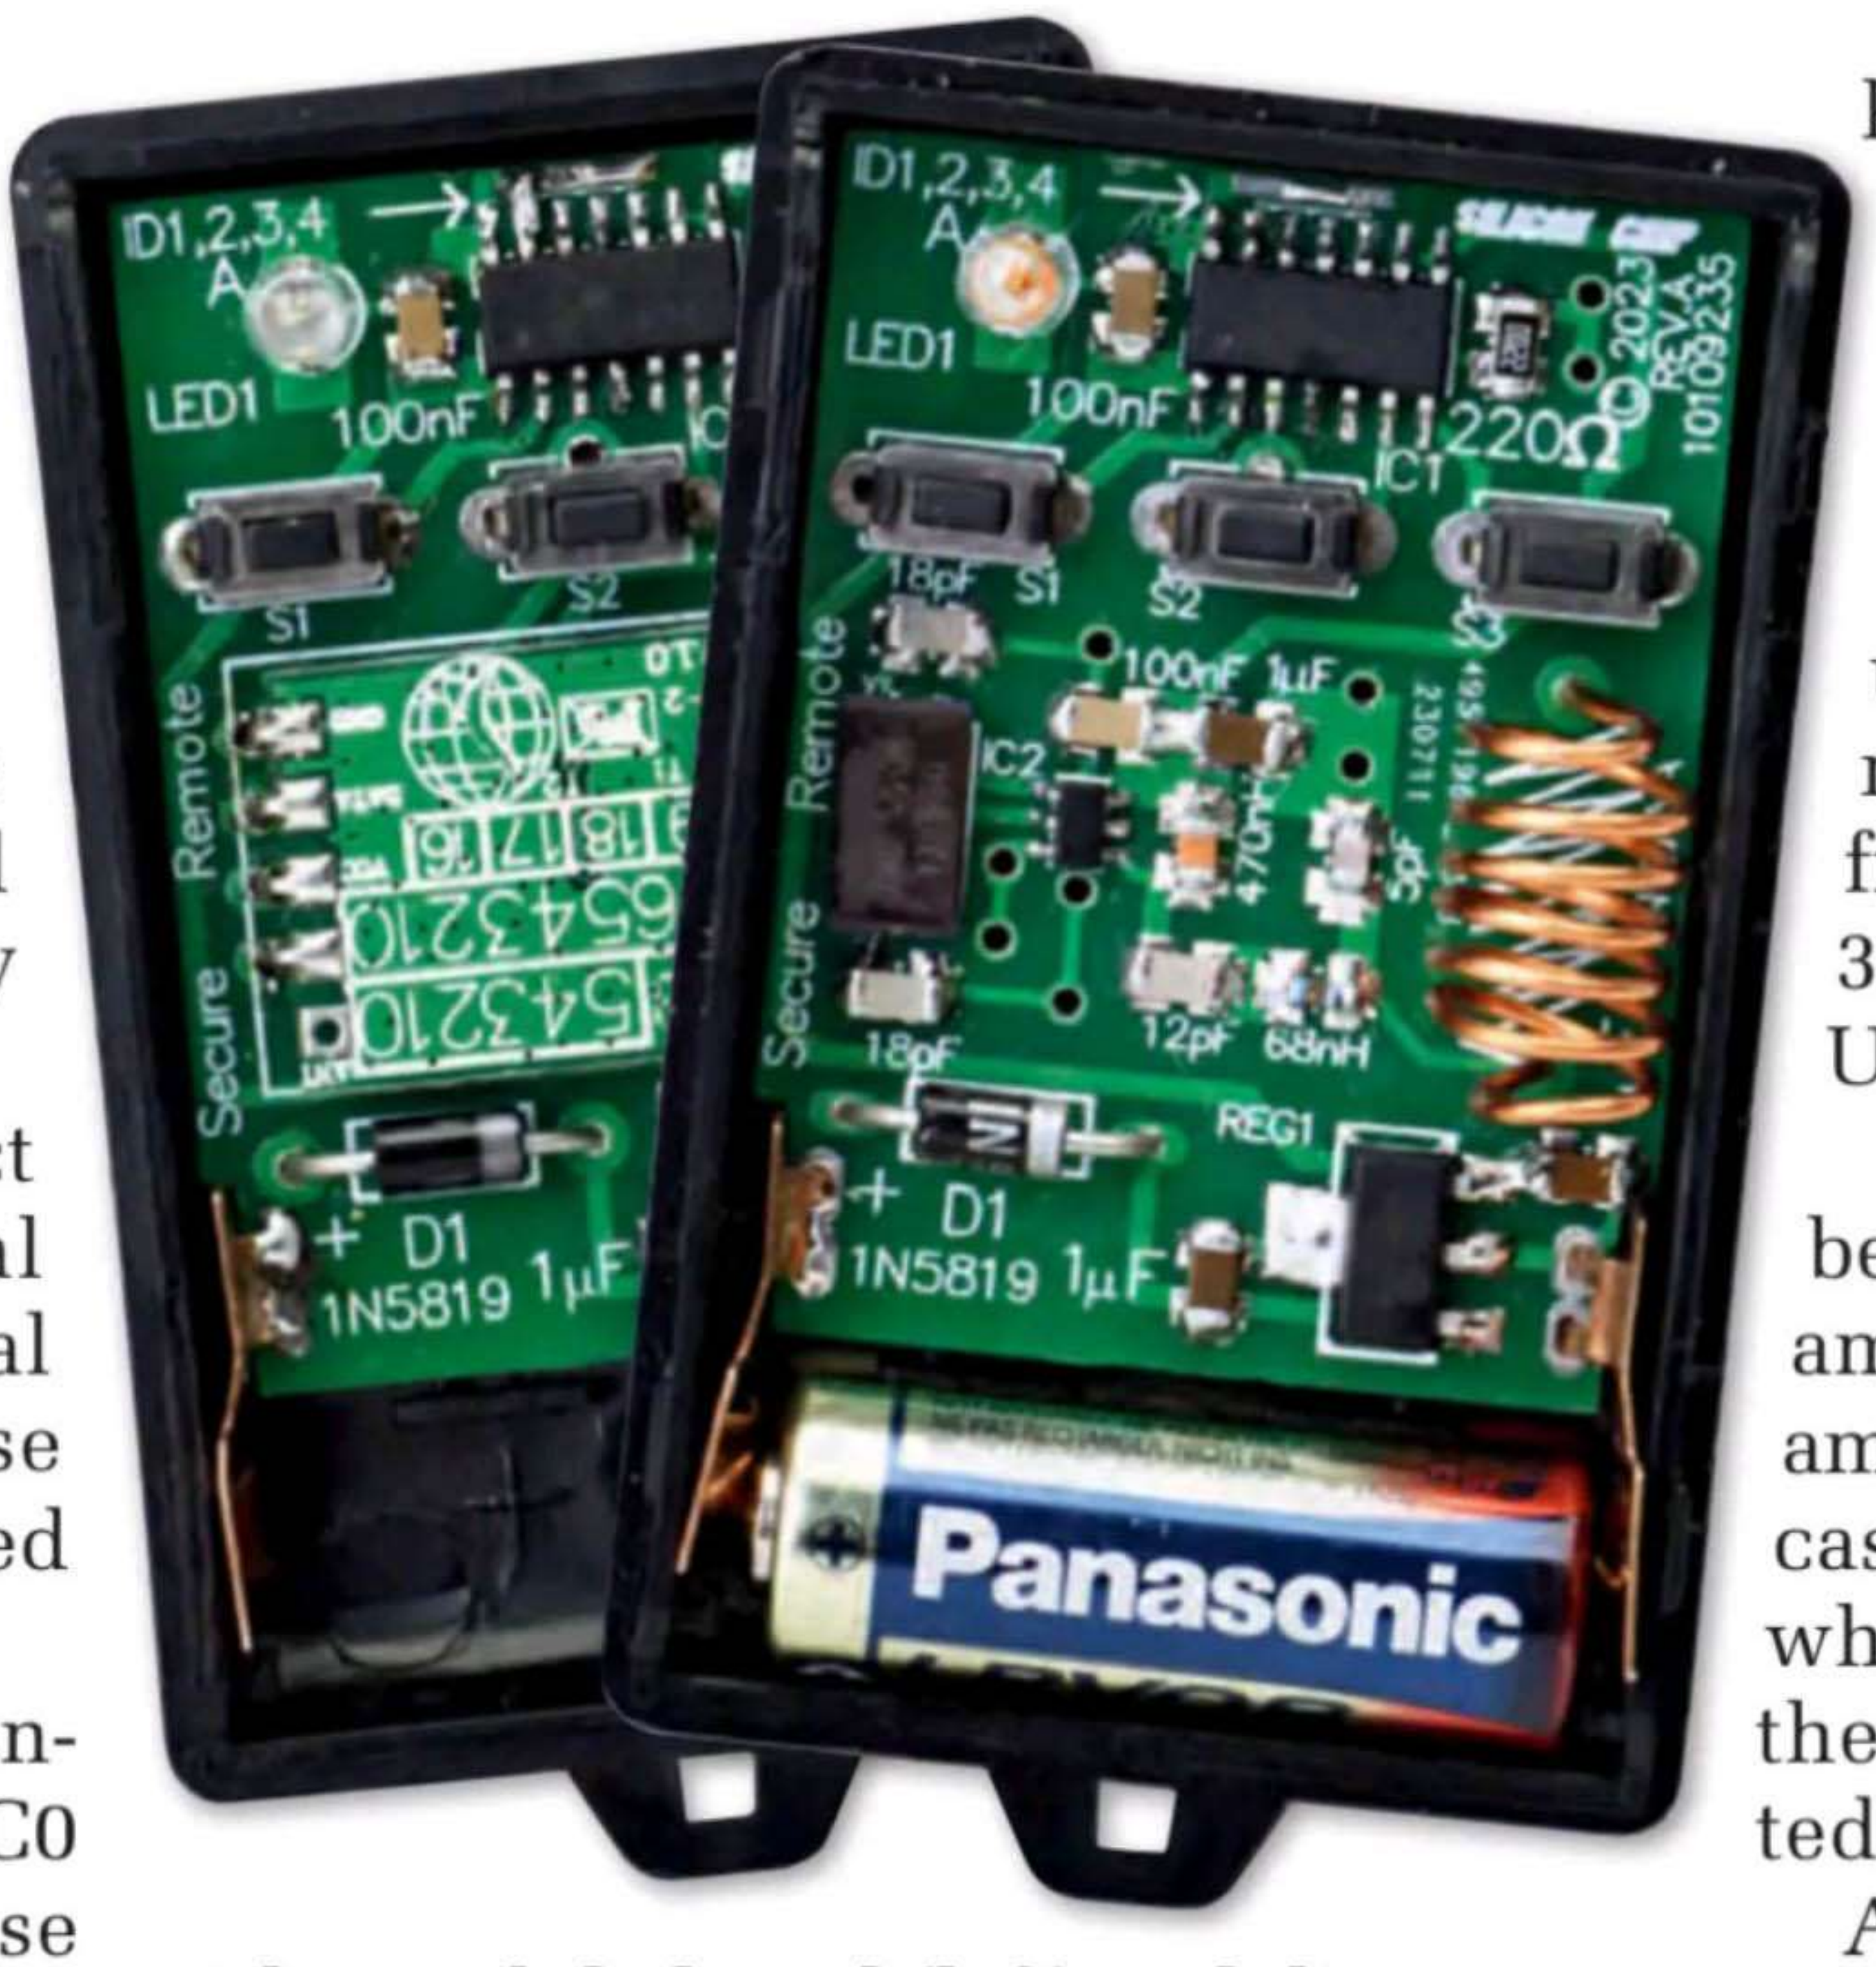
D1,2,3,4
LED1
100nF
S1
S2
S3
Remote
Secure
D1
1N5819
1µF
IC1
IC2
220Ω
18pF
100nF
470nF
12pF
68nH
5pF
REG1
Panasonic

The module-based (left) and discrete (right) versions of the transmitter PCB shown enlarged. We have used an A23 12V battery, which fits snugly with the recommended battery clips.

The pullups for pushbutton switches S1-S3 are left on permanently since they are only pressed momentarily. IC1 is programmed to wake up from its sleep condition when any one of switches S1-S3 is pressed and the corresponding input goes low. It then runs the program to send the rolling code for the function associated with the pressed switch.

When a button is pressed, the micro drives its RC2 and RC1 digital outputs high, to 3.3V. These are connected in parallel to power the UHF transmitter (module or discrete components). This way, UHF transmit circuitry only draws current from the battery when it is in use.

With the transmitter powered up, IC1 sends the rolling code and registration codes on the data line from its digital output RA4 (pin 3). This feeds the data input of the UHF circuitry.

UHF code transmission switches between two different carrier wave amplitudes, a technique known as amplitude shift keying (ASK). In this case, there is no UHF transmission when the digital signal is low, but the 433.92MHz carrier is transmitted when the digital signal is high.

After sending the code, IC1 powers down the UHF transmitter and returns to sleep mode.

Discrete UHF circuitry

Referring to the additional UHF transmission circuitry in Fig.2, the MICRF113 is a single-chip ASK UHF transmitter IC. Its transmission frequency is set using a crystal oscillator multiplied by 32 within IC2 to produce the UHF carrier. So the 13.56MHz crystal results in a 433.92MHz carrier.

This matches the carrier frequency used by most UHF ASK transmitter/receiver modules available for low-power UHF data transmission.

IC2's power rail at pin 3 is bypassed with 100nF & 1µF ceramic capacitors

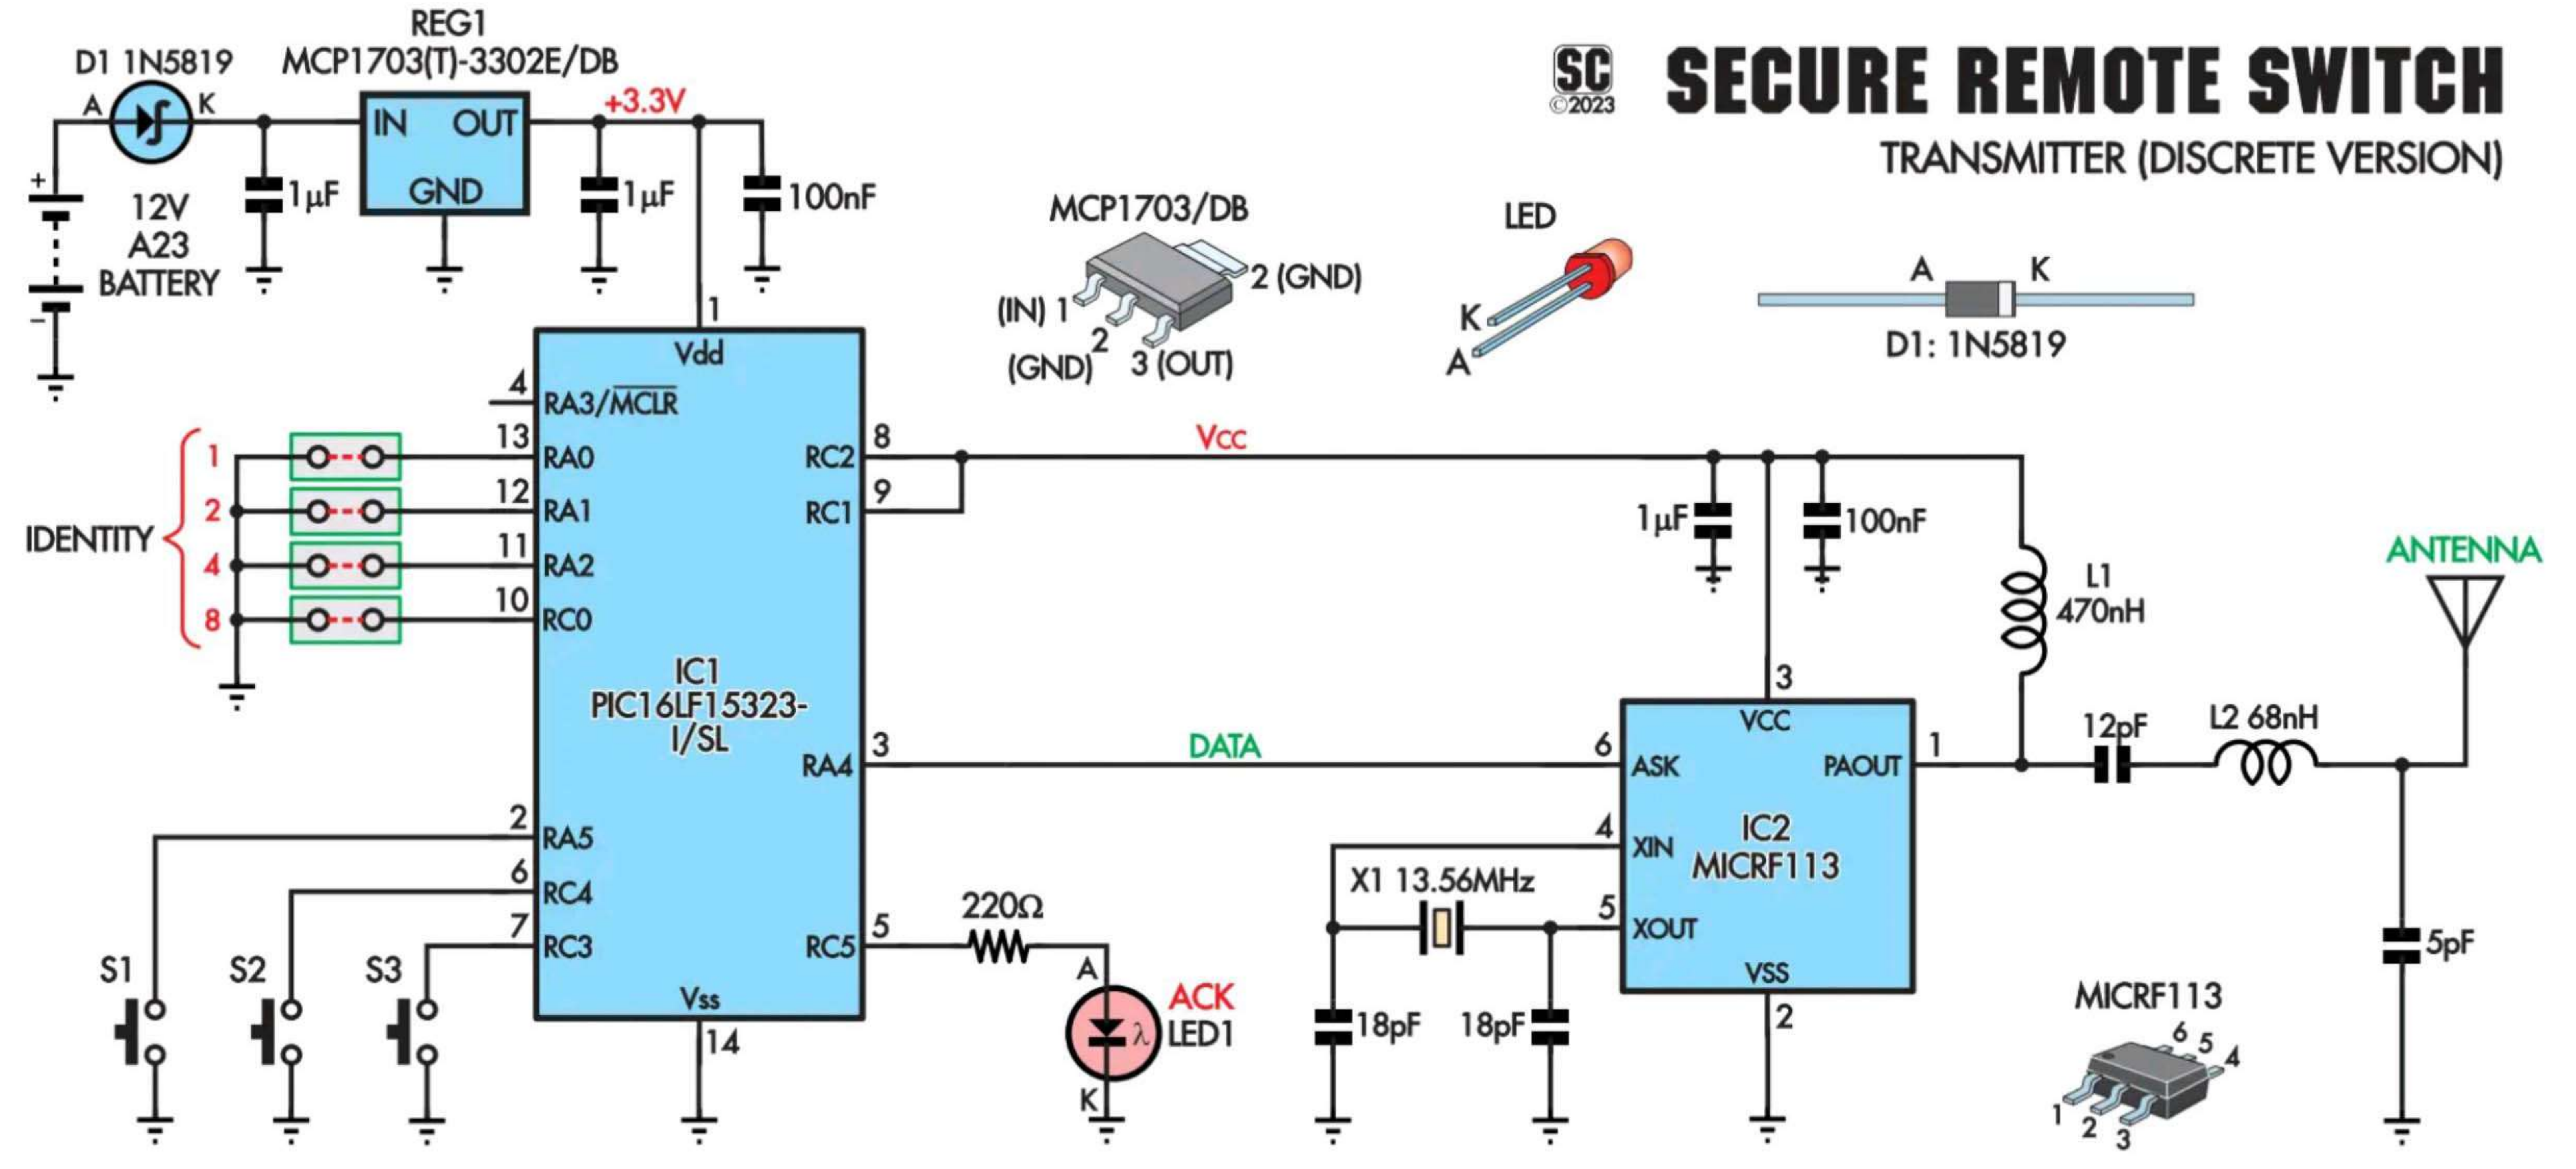
SC ©2023
SECURE REMOTE SWITCH
TRANSMITTER (DISCRETE VERSION)
D1 1N5819
A
K
REG1
MCP1703(T)-3302E/DB
IN
OUT
GND
+3.3V
12V A23 BATTERY
1µF
1µF
100nF
MCP1703/DB
(IN) 1
2 (GND)
(GND) 2
3 (OUT)
LED
K
A
D1: 1N5819
1
Vdd
4
RA3/MCLR
13
RA0
12
RA1
11
RA2
10
RC0
IDENTITY
1
2
4
8
IC1
PIC16LF15323-I/SL
RC2
8
RC1
9
Vcc
RA4
3
DATA
2
RA5
6
RC4
7
RC3
RC5
5
Vss
14
S1
S2
S3
220Ω
ACK
LED1
A
K
1µF
100nF
L1
470nH
ANTENNA
3
VCC
6
ASK
PAOUT
1
4
XIN
IC2
MICRF113
5
XOUT
VSS
2
X1 13.56MHz
18pF
18pF
12pF
L2 68nH
5pF
MICRF113
1 2 3 4 5 6

Fig.2: the left side of the discrete version of the transmitter circuit is identical to Fig.1. This time, the MICRF113 ASK IC generates a 433.9MHz carrier from the 13.56MHz crystal and switches it on and off based on the digital signal at its ASK input (pin 6). Inductor L1 is its output load, while L2 and the 12pF & 5pF capacitors filter out unwanted harmonics.

Parts List – Secure Remote Switch (Receiver)

1 double-sided plated-through PCB coded 10109231, 70 × 96.5mm
1 set of front and rear panel labels
1 Ritec 105 × 80 × 33mm plastic enclosure **[Altronics H0191]**
1 433.9MHz UHF ASK receiver (RX1) **[rfsolutions AM-RX12A-433P]**
1 10A SPDT relay (12V or 24V coil) (RLY1)
1 subminiature SPDT PCB-mount momentary horizontal pushbutton switch (S1)
1 button cap for S1
2 SPST PCB-mount tactile micro switches (S2, S3)
1 4-bit (0-9 & A-F) 6-pin BCD PCB-mount rotary switch (S4)
1 subminiature SPDT PCB-mount horizontal toggle switch (S5)
1 PCB-mount barrel socket, 2.1mm or 2.5mm inner diameter (CON1)
1 2-way screw terminal, 5/5.08mm pitch (CON2)
1 3-way screw terminal, 5/5.08mm pitch (CON3)
1 10kΩ miniature single-turn top-adjust trimpot (code 103) (VR1)
3 2-way pin headers, 2.54mm pitch (JP1-JP3)
3 jumper shunts (JP1-JP3)
1 20-pin DIL IC socket (for IC1)
1 PG7 (3-6.5mm cable) or PG9 (4-8mm cable) cable gland for rear panel
1 169mm length of 0.8mm diameter enamelled copper wire
1 169mm length of 1mm diameter heatshrink tubing (**optional**)

Semiconductors
1 PIC16F1459-I/P programmed with 1010923R.HEX, DIP-20 (IC1)
1 7805 5V 1A linear regulator, TO-220 (REG1)
1 BC337 500mA NPN transistor, TO-92 (Q1)
2 1N4004 400V 1A diodes, DO-41 (D1, D2)
1 3mm high-brightness red LED (LED1)
2 5mm high-brightness LEDs (eg, red & green) (LED2, LED3)

Capacitors
1 100µF 25V PC electrolytic •
1 100µF 16V PC electrolytic
1 10µF 35V PC electrolytic •
2 100nF MKT polyester or ceramic (code 104 or 100n)
• *can be 16V rated for 12V supply*

Resistors *(all ¼W, 1% metal film unless noted)*
5 10kΩ 3 560Ω 1 330Ω
470Ω 1W for 24V supply, 100Ω ½W for 12V supply (R1)

DC IN @100mA
12V
24V
-⊖-◉-⊕-+

The rear of the receiver case includes the power socket and cable glands for wiring to the relay terminals.

while the supply current for IC2's RF output stage is via a 470nH inductor acting as a driver load. The following 12pF series capacitor and 68nH inductor plus the 5pF capacitor to ground act as a filter to remove second and third harmonics from the UHF signal before it passes to the antenna.

Any inductor used for the output stage and filter circuit must have a self-resonance (SR) frequency above 433.92MHz; otherwise, it will not function as an inductor at that frequency. This is a critical requirement for any substitute components to those specified in the parts list.

Power supply

In both cases, IC1 is powered using an A23 12V battery and a 3.3V low-quiescent-current low-dropout voltage regulator (REG1). This supplies the UHF transmitter section as well as the microcontroller. REG1 typically draws a 2µA quiescent current at 25°C, although that could be as high as 5µA over the range of -40°C to +125°C.

With IC1 in sleep mode, it draws a typical standby current of 60nA from its 3.3V supply and so can essentially be ignored compared to the regulator's quiescent current.

We measured the quiescent current draw from the 12V battery on our two prototypes at 2.7µA and 3µA, respectively. When a switch is pressed on the transmitter, that increases but only briefly, so that does not affect the long-term battery life much.

During transmission, the current draw from the battery briefly rises to about 10mA. If you keep holding one of the buttons down after the transmission is complete, the current will drop to about 220µA until the button is released. This is due to the pushbutton switch pullup current.

Considering the low quiescent current and intermittent bursts of higher current when transmitting, battery life should be more than two years with typical use.

Receiver circuit

The receiver circuit (Fig.3) uses a PIC16F1459-I/P microcontroller (IC1) and UHF receiver module with an onboard wire antenna to provide a good reception range.

When no signal is present, the receiver's output produces random noise since the module's automatic gain control (AGC) is at its maximum. Upon reception of a 433.92MHz signal, the receiver gain is reduced for best reception without overload, and the coded signal from the data output of the module is delivered to the RC7 digital input of IC1 (pin 9).

IC1 flashes the Acknowledge LED (LED2) whenever a valid signal is received. This also doubles as a relay-on indicator. It is lit when the relay is on and off when the relay is off.

The RC5 digital output of IC1 (pin 5) drives NPN transistor Q1, which switches the relay coil. When RC5 goes high, it delivers current to transistor Q1's base, and Q1 powers RLY1. Diode D2 clamps the back-EMF that causes a voltage spike at the collector of Q1 as the relay switches off. The relay contacts are rated at 10A for AC or DC.

The unit can be set up to power the relay for a fixed period when a transmitter button is pressed (or S1 on the receiver) or toggle it on or off for each button press. This on/off functionality can be set differently for the transmitter buttons and the onboard pushbutton, S1. Since the transmitters have three buttons, they can provide different functions (more on that shortly).

When jumper JP3 is closed, the relay switches on with one press of onboard button S1 and off with the next. When JP3 is open, the relay is switched on for a fixed time with a press of S1 and switches off automatically at the end of this period – see Table 3.

The remote control has three buttons; usually, S1 on the remote switches the relay on, and it is then switched off with the timer. S2 switches it on continuously (or for a much longer time if JP2 is inserted), and S3 switches it off – see Table 2.

The timer period is set using trimpot VR1. The trimpot wiper can be adjusted from 0V through to 5V; this voltage is monitored at the AN6 analog input of IC1, which converts the voltage into setting a period from 0.25 seconds to 60 seconds or one minute to four hours and 30 minutes, depending on the settings of JP1 & JP2 (see Tables 1 & 2).

Table 1 – JP1 timer settings

JP1	Timer range
Out	0.25-60s (1x)
In	1m-4.5h (255x)

Table 2 – JP2 settings

TX button	Function with JP2 out	Function with JP2 in
S1	Relay on with a timer, range per JP1	Relay on with a timer, 0.25-60s
S2	Relay on continuously	Relay on with a timer, 1m-4.5h
S3	Relay off	Relay off

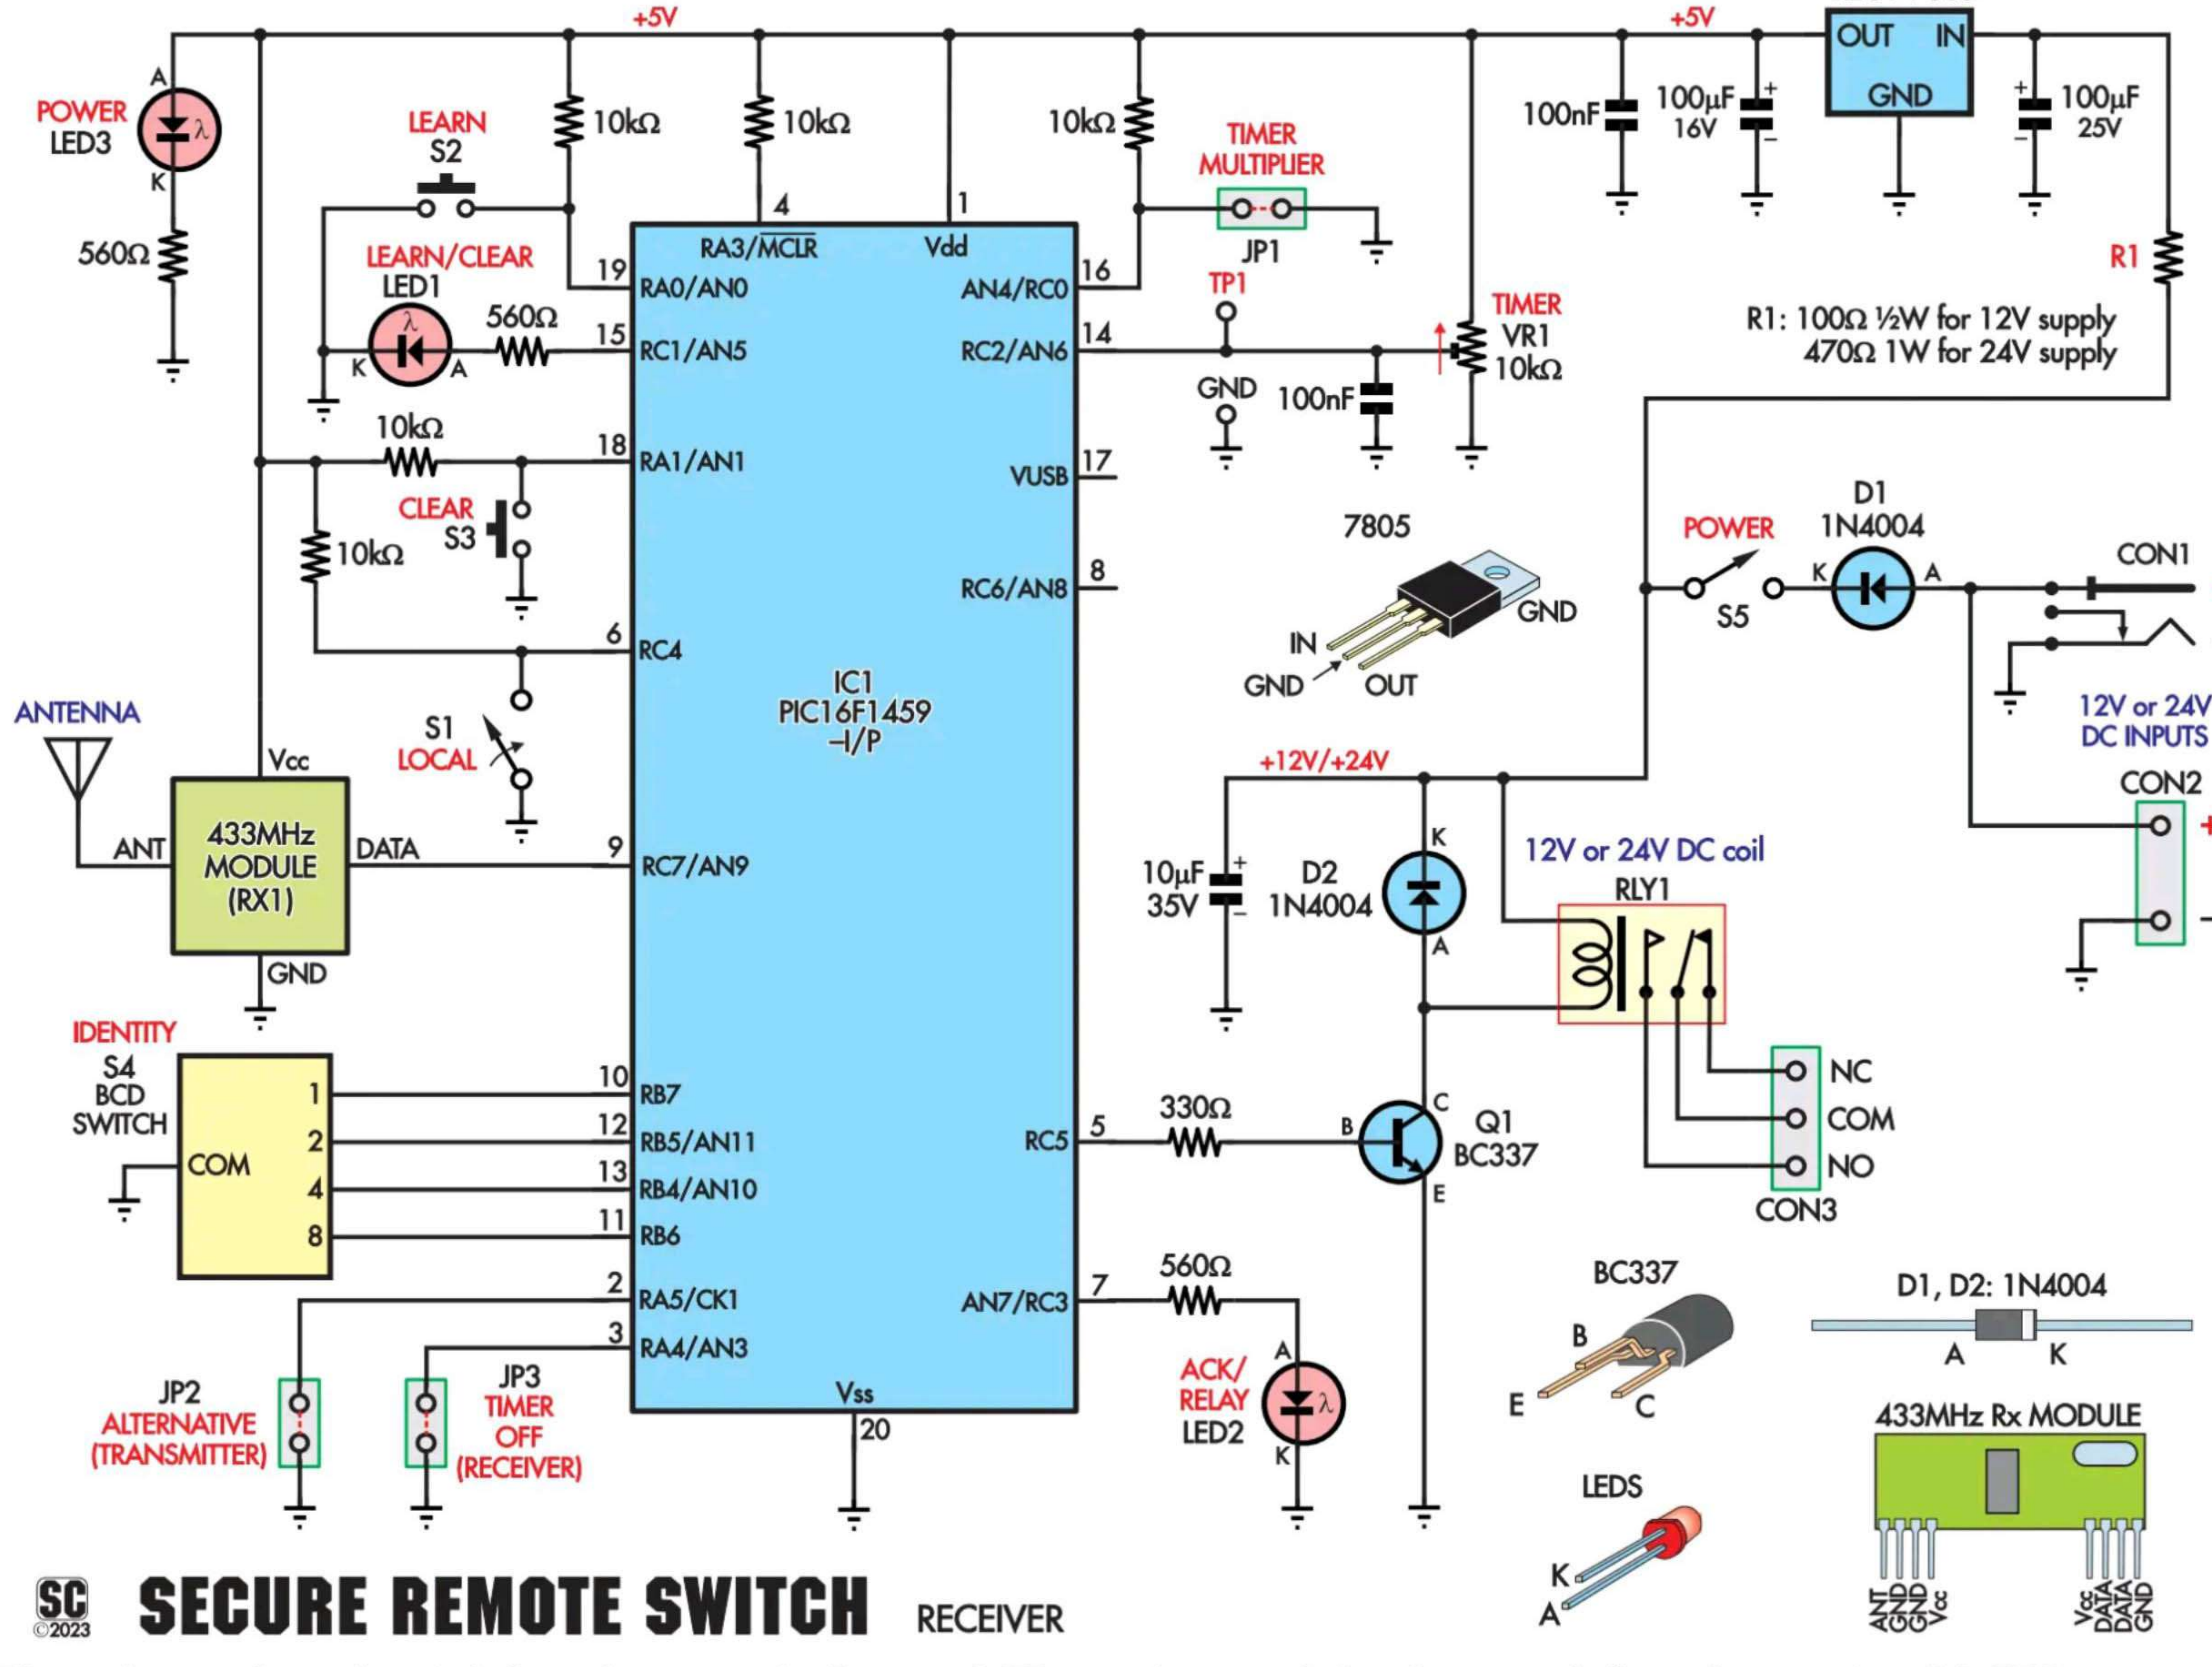
+5V
POWER
LED3
A
K
560Ω
LEARN
S2
10kΩ
10kΩ
10kΩ
TIMER
MULTIPLIER
JP1
LEARN/CLEAR
LED1
560Ω
19
RA0/AN0
15
RC1/AN5
4
RA3/MCLR
1
Vdd
16
AN4/RC0
14
RC2/AN6
TP1
GND
100nF
TIMER
VR1
10kΩ
10kΩ
18
RA1/AN1
CLEAR
S3
10kΩ
6
RC4
S1
LOCAL
17
VUSB
8
RC6/AN8
IC1
PIC16F1459
–I/P
ANTENNA
ANT
433MHz
MODULE
(RX1)
Vcc
DATA
GND
9
RC7/AN9
IDENTITY
S4
BCD
SWITCH
COM
1
2
4
8
10
RB7
12
RB5/AN11
13
RB4/AN10
11
RB6
2
RA5/CK1
3
RA4/AN3
Vss
20
5
RC5
7
AN7/RC3
JP2
ALTERNATIVE
(TRANSMITTER)
JP3
TIMER
OFF
(RECEIVER)
330Ω
560Ω
ACK/
RELAY
LED2
Q1
BC337
B
C
E
REG1 7805
OUT
IN
GND
100nF
100µF
16V
100µF
25V
R1
R1: 100Ω ½W for 12V supply
470Ω 1W for 24V supply
7805
IN
GND
OUT
POWER
S5
D1
1N4004
CON1
+
–
12V or 24V
DC INPUTS
CON2
+12V/+24V
10µF
35V
D2
1N4004
12V or 24V DC coil
RLY1
NC
COM
NO
CON3
BC337
D1, D2: 1N4004
433MHz Rx MODULE
ANT
GND
GND
Vcc
Vcc
DATA
DATA
GND
LEDS
SC
©2023
SECURE REMOTE SWITCH
RECEIVER

Fig.3: the receiver circuit is based on a prebuilt 433.9MHz receiver module, shown at left, and a 20-pin 8-bit PIC microcontroller, IC1. When IC1 receives a valid rolling code, it brings its pin 5 high to power NPN transistor Q1 which switches the relay coil. The relay is a 12V or 24V DC coil type to match the supply voltage.

Rolling code transmission format

The rolling code is transmitted using UHF ASK in Manchester code. A zero-bit is sent as a 512μs period of no transmission followed by a 512μs burst of 433.9MHz carrier. In contrast, a one-bit is transmitted as a 512μs burst of 433.9MHz carrier followed by a 512μs period of no signal.

Each transmission consists of four start bits, an eight-bit identifier, a 48-bit code and four stop bits, for a total of 64 bits. The start bits include a 16.4ms gap between the second and third start bit, while the code scramble value is altered on each transmission with 32 variations.

Unique codes are generated with a 48-bit seed, 24-bit multiplier & 8-bit increment value. That is initially set by a unique identifier within IC1 on the transmitter.

The registration code is sent as two blocks. Block 1 sends four start bits, the eight-bit identifier, a 32-bit seed code and four stop bits. Block 2 sends four start bits, the 24-bit multiplier, the eight-bit increment and eight-bit scramble values and four stop bits. Again, the start bits include a 16.4ms gap between the second and third start bit.

Table 3 – JP3 settings

JP3	Onboard S1 function
Out	Off if already on, otherwise on for a time set by JP1 and VR1 (see Table 1)
In	Toggle on/off

Table 4 – period vs TP1 voltage

TP1	Time with JP1 out	Time with JP1 in
0V	0.25s	1m
1.25V	15s	1h 7.5m
2.5V	30s	2h 15m
3.75V	45s	3h 22.5m
5V	60s	4h 30m

IC1's digital input RC0 for JP1 has an external 10kΩ pullup resistor. If JP1 is inserted, this pin is held low. IC1 senses that and, in that case, changes the maximum timer setting from one minute to 4 hours and 30 minutes.

You can monitor the timer setting voltage between test points TP1 and GND. Table 4 shows the typical periods for five different voltages in each range.

Transmitter Identity

The receiver Identity selection is made using a BCD rotary switch (S4) with 16 positions, labelled 0-9 and then A-F. Those hexadecimal values correspond to 0-15 in decimal, with A-F representing 10-15. This switch is only monitored by IC1 for lockout selections; it plays no part in the keyfob transmitter registration.

S4's four contacts connect to the RB7, RB6, RB5 and RB4 digital inputs of IC1. These all have internal pullups, so the inputs are at 5V when the corresponding switch is not closed. All four inputs are high when the BCD switch is set at 0. Position 1 on the switch has the '1' output at RB7 pulled low, while position 15 (or F) sets all four pins to 0V.

Deregistration & registration

S3 is used for deregistering a transmitter. Pressing S3 for more than one second will deregister the transmitter specified by the BCD switch, preventing it from operating the receiver again. Successful deregistration will be acknowledged by the Learn/Clear LED (LED1) lighting. Table 5 shows the identity selection coding for both the transmitter and receiver.

The Learn switch (S2) tells the program within IC1 to be ready to accept the synchronising signal from a handheld remote. The Learn/Clear LED (LED1) stays lit while waiting for a signal from the remote unit. It extinguishes once the synchronising signal has been correctly received.

Power supply

The receiver can be powered from 12V or 24V DC, from a DC plugpack or similar DC supply; some garage door controllers have DC supply terminals that could also be used. Regardless of the source, power can be connected via CON1 (a barrel socket) or two-way screw terminal CON2.

Reverse polarity protection is via diode D1, which only allows current to flow into the circuit if the supply polarity is correct.

The relay has a 12V or 24V DC coil, matching supply voltage. For 24V, a 470Ω 1W resistor (R1) reduces the voltage applied to 5V regulator REG1. For a 12V DC input, a 100Ω ½W resistor is used instead.

The 470Ω resistor reduces the dissipation in REG1 when the supply is at 24V. This resistor also filters the DC supply to REG1 in conjunction with the 100μF input capacitor, removing most of the noise from a switchmode supply that could otherwise affect the UHF receiver sensitivity.

For a 24V DC supply, the 100μF capacitor is rated at 25V, and the 10μF capacitor bypassing the relay supply is 35V. For a 12V supply, the capacitors can all be rated at 16V.

Table 5 – Transmitter Identity selection

Receiver	Transmitter	Transmitter	Transmitter	Transmitter
S4	'1'	'2'	'4'	'8'
0	open	open	open	open
1	shorted	open	open	open
2	open	shorted	open	open
3	shorted	shorted	open	open
4	open	open	shorted	open
5	shorted	open	shorted	open
6	open	shorted	shorted	open
7	shorted	shorted	shorted	open
8	open	open	open	closed
9	shorted	open	open	closed
A (10)	open	shorted	open	shorted
B (11)	shorted	shorted	open	shorted
C (12)	open	open	shorted	shorted
D (13)	shorted	open	shorted	shorted
E (14)	open	shorted	shorted	shorted
F (15)	shorted	shorted	shorted	shorted

Practical Electronics

UK readers SAVE £1 on every issue

SUBSCRIBE NOW!

01202 087631

Take out a one-year subscription and save more than £10 over the year.
Even better – save over £1 per issue if you subscribe for two years – a total saving of £26.

Overseas rates represent exceptional value

You also:
- **Avoid any cover price increase for the duration of your subscription**
- **Get your magazine delivered to your door each month**
- **Guarantee a copy, even if your local newsagent sells out**

Order by:
- **Phone or post with a cheque, postal order or credit card**
- **Or order online at: www.electronpublishing.com**

Practical Electronics

SUBSCRIPTION PRICES

Subscriptions for delivery direct to any address in:

UK: 6-months £33.99, 12-months £59.99, 24-months £114.99

Europe Airmail: 6-months £38.99, 12-months £69.99, 24-months £129.99

Rest Of The World Airmail: 6-months £46.99, 12-months £84.99, 24-months £164.99

Cheques (in **£ sterling only**) payable to *Practical Electronics* and sent to: Practical Electronics Subscriptions, PO Box 6337, Bournemouth BH1 9EH, United Kingdom

Tel: 01202 087631
Email: pesubs@selectps.com

Also via our online shop at**: www.electronpublishing.com**

Subscriptions start with the next available issue. We accept MasterCard or Visa.
(For past issues see the Back Issues page.)

DIGITAL SUBSCRIPTIONS

Electronic subscriptions are available from £29.99 for 12 months, more details at: **www.electronpublishing.com**

SUBSCRIPTION ORDER FORM

☐ 6 Months: UK £33.99, Europe £38.99 (Airmail), Rest of the World £46.99 (Airmail)
☐ 1 Year: UK £59.99, Europe £69.99 (Airmail), Rest of the World £84.99 (Airmail)
☐ 2 Years: UK £114.99, Europe £129.99 (Airmail), Rest of the World £164.99 (Airmail)

To: **Practical Electronics Subscriptions**
PO Box 6337
Bournemouth BH1 9EH
United Kingdom

Tel: 01202 087631
Email: pesubs@selectps.com

I enclose payment of £ ....
(cheque/PO in £ sterling only)
payable to *Practical Electronics*

☐ Please charge my Visa/Mastercard

My card number is: ....
Please print clearly, and check that you have the number correct

Card Security Code .... Valid From Date.......................
(The last 3 digits on or just under the signature strip)

Card Ex. Date ....

Name ....

Address ....

Post code .... Tel. ....

Email ....

Subscriptions start with the next available issue.

Test equipment review

Using Electronic Modules with Jim Rowe

ZPB30A1 Module

- 60W Programmable DC Load

- Battery Capacity Tester

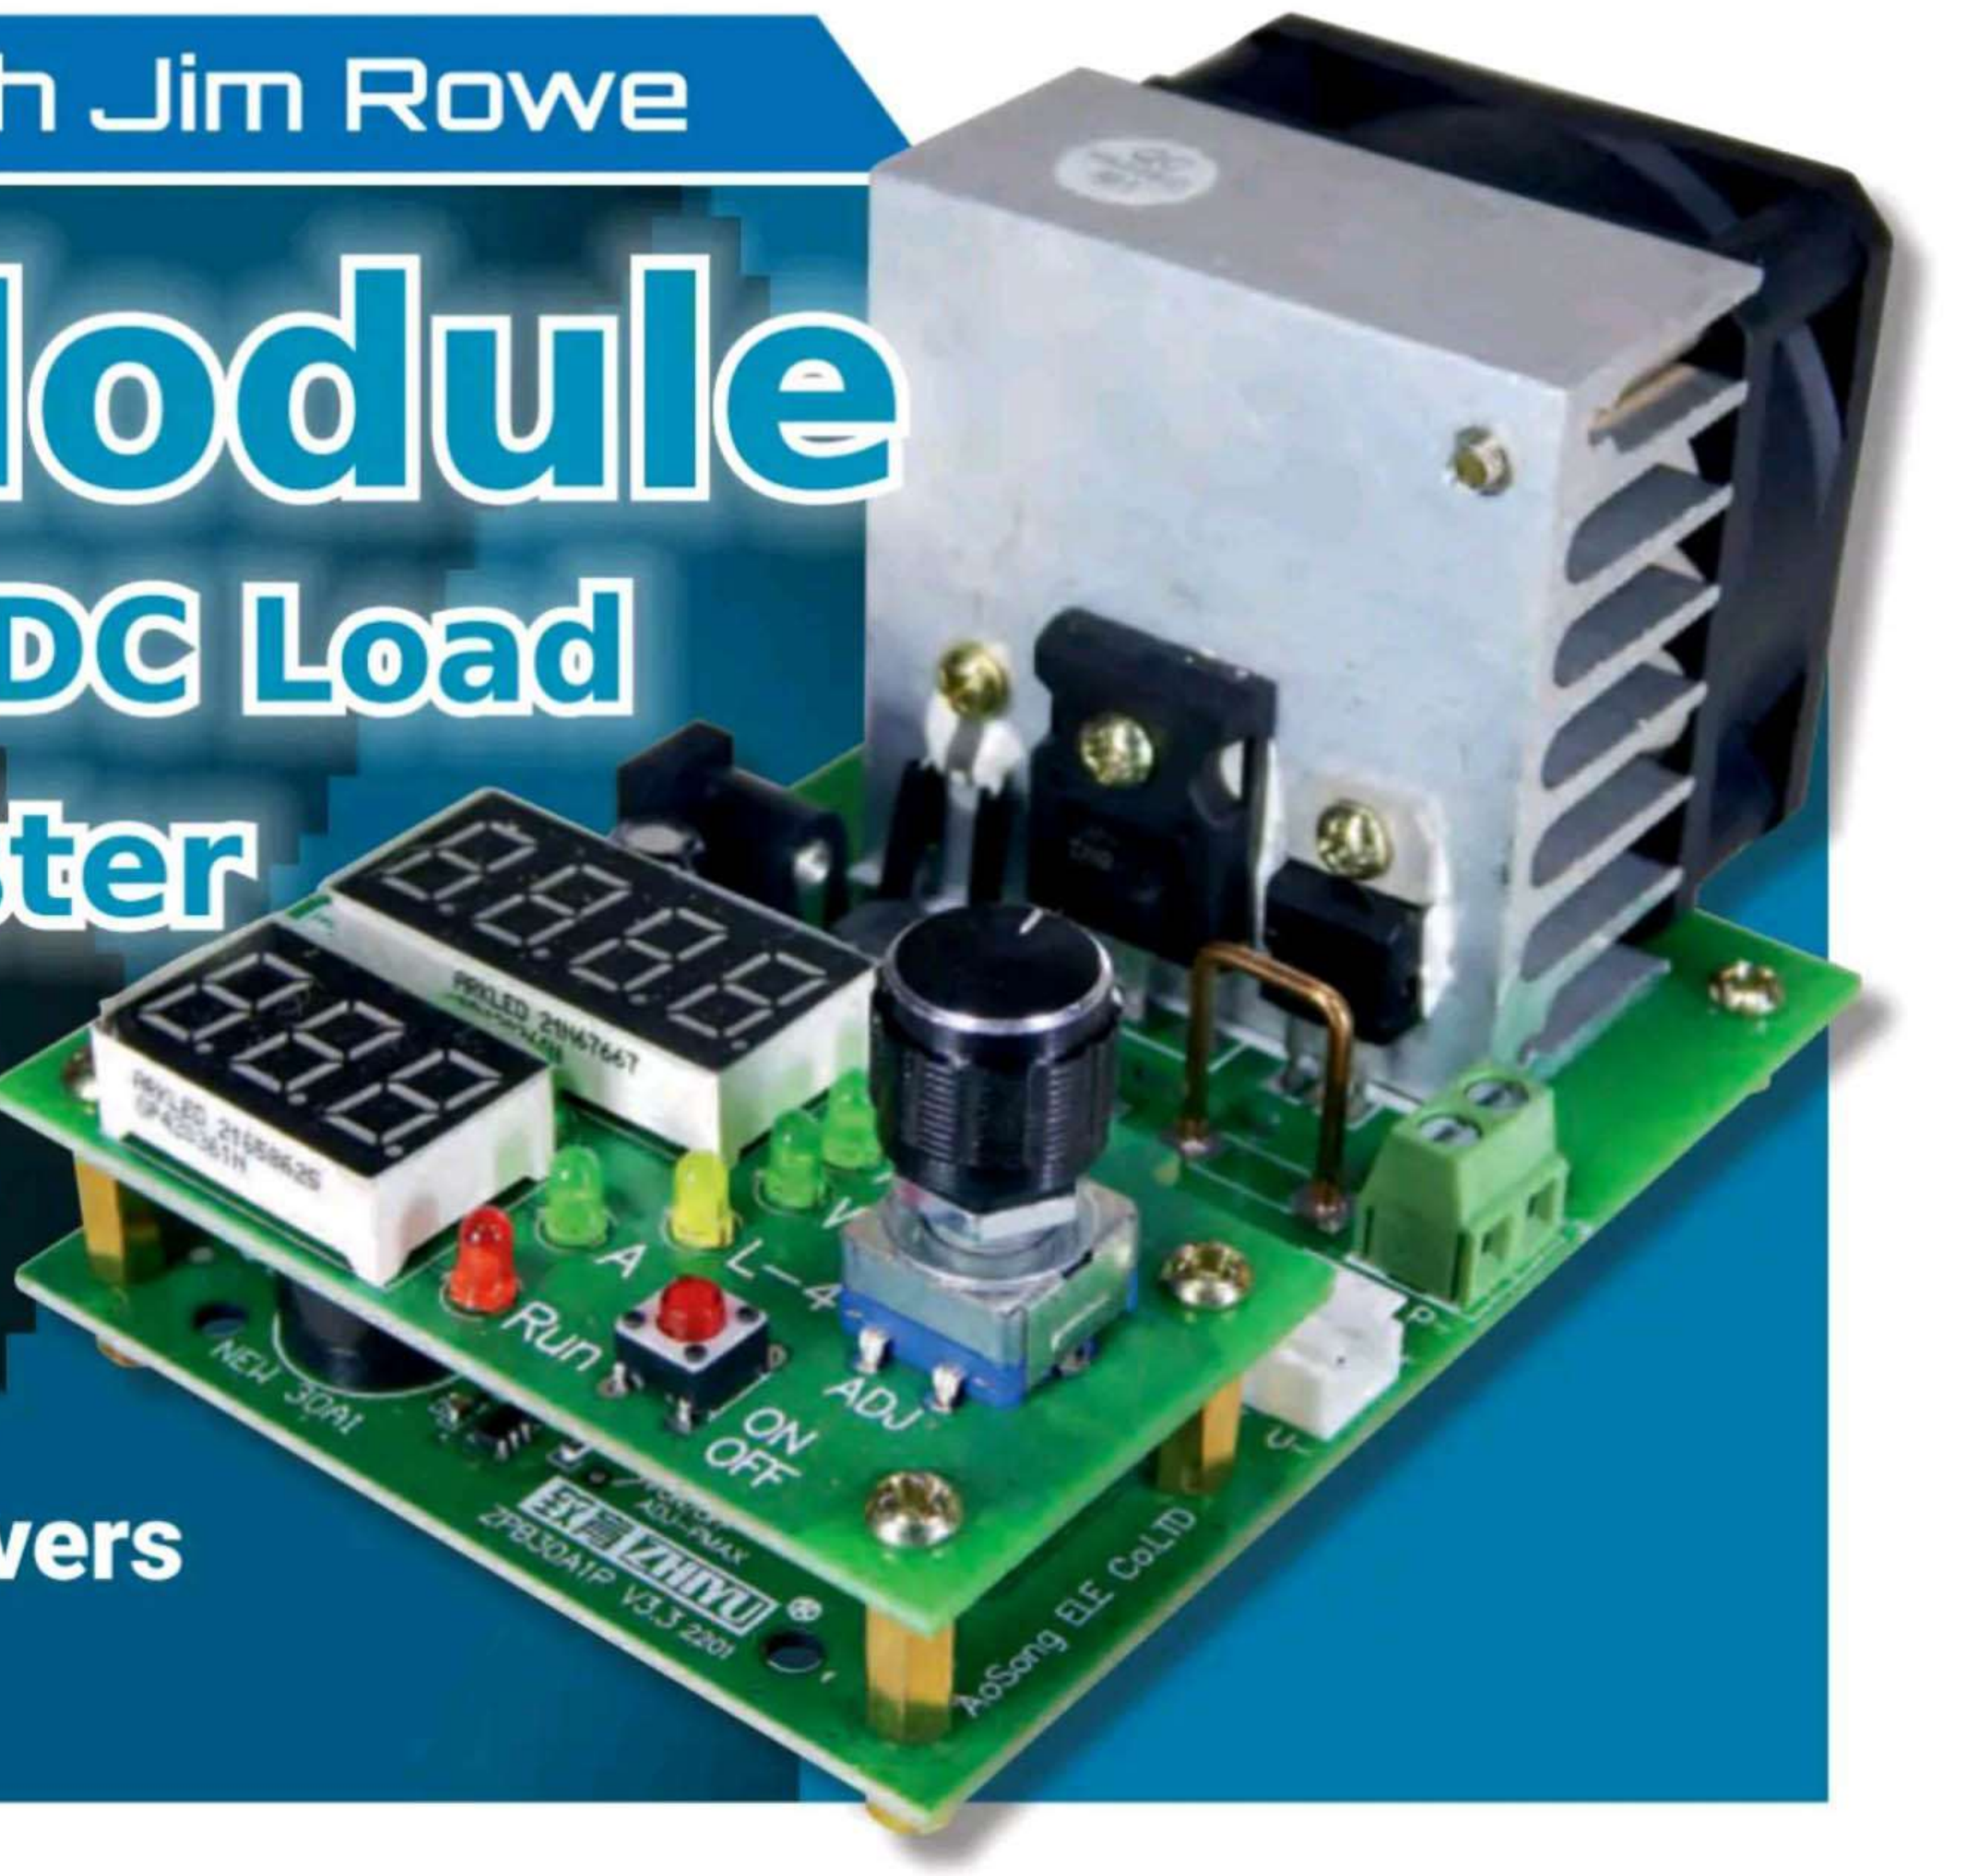

This programmable constant-current DC load can be used for testing power supplies or checking the capacity of storage batteries. It is essentially self-contained and delivers good value for the money.

The ZPB30A1 module carries the brand name Zhiyu, but it seems to be made in China by a firm called AoSong ELE Co Ltd. As you can see from the photos, it has two PCBs, with the smaller one (69 × 36mm) mounted above the larger main PCB that measures 100 × 69mm.

The 50 × 50 × 23mm heatsink is at the rear of the larger PCB, along with its associated cooling fan on the finned side, for cooling the main load transistor. The fan extends past the rear of the main PCB by about 11mm.

Mounted on the flat front of the heatsink is the power Mosfet that acts as the controlled load (in the centre), with a thermistor to its left used for sensing its temperature. To its right is a dual schottky diode that protects the power transistor from damage due to reversed voltage polarity.

The smaller PCB is the control and display board, or the main 'user interface'. Its four-digit 7-segment LED displays the voltage, battery capacity or various control and error messages. In contrast, the three-digit LED display below it is mainly used to show the current flow.

Six additional green 3mm LEDs indicate which parameter is being displayed or which 7-segment digit is being adjusted, while a red 3mm LED indicates when the module is running. The function of the 3mm yellow LED function is not explained; it is labelled "L-4" and seems to be a recent addition to the latest (V3.3) version of the module.

On the right of the smaller PCB are the two controls. The first is a rotary encoder, which changes modes and adjusts current and voltage values. The second is a small pushbutton used to confirm the module's current and termination voltage settings and as an on/off control.

Currently, the ZPB30A1 module is available from several sources on the internet, including Banggood (www.banggood.com/search/1146280.html) and many suppliers on AliExpress and eBay, ranging from £7.65 plus £3.78 for delivery to £19.25 plus £2.30 for delivery.

I ordered one from Banggood at a price near the high end, and after the usual wait, it arrived safely – even though it was only wrapped in bubble wrap inside a plastic bag.

What it does

The module has two modes of operation. One is to serve as a programmable constant-current DC load, while the other is to test the capacity of storage batteries like Li-ion, Nicad or lead-acid batteries. The basic specifications are shown in the sidebar.

The module's internal circuitry is

FEATURES & SPECIFICATIONS

- Test modes: programmable constant-current DC load ("Fun1") or battery capacity tester ("Fun2")
- Maximum dissipation: 60W
- Operating voltage range: 1-30V (separate 12V 500mA supply required)
- Operating current range: 0.1-9.99A in steps of 0.01A (10mA)
- Rated current measurement accuracy: ±(0.7% + 10mA)
- Test termination voltage range: 1-25V
- Voltage measurement: directly at the P+ and P– terminals or remotely for four-terminal measurements
- Voltage measurement accuracy: ±(1% + 0.02V)
- Battery capacity maximum values: 999.9Ah, 9999Wh
- Battery capacity test accuracy: 2.5% @ 0.5A, 1.5% @ 2A or 1.2% @ 5A+
- Protection: over-temperature ("otP"), transient over-power ("oPP"), over-voltage ("ouP"), reverse polarity ("Err3") and abnormal voltage ("Err6")
- Fan control: automatic, temperature-controlled
- Size: 69 × 111 × 57mm
- Weight: 270g

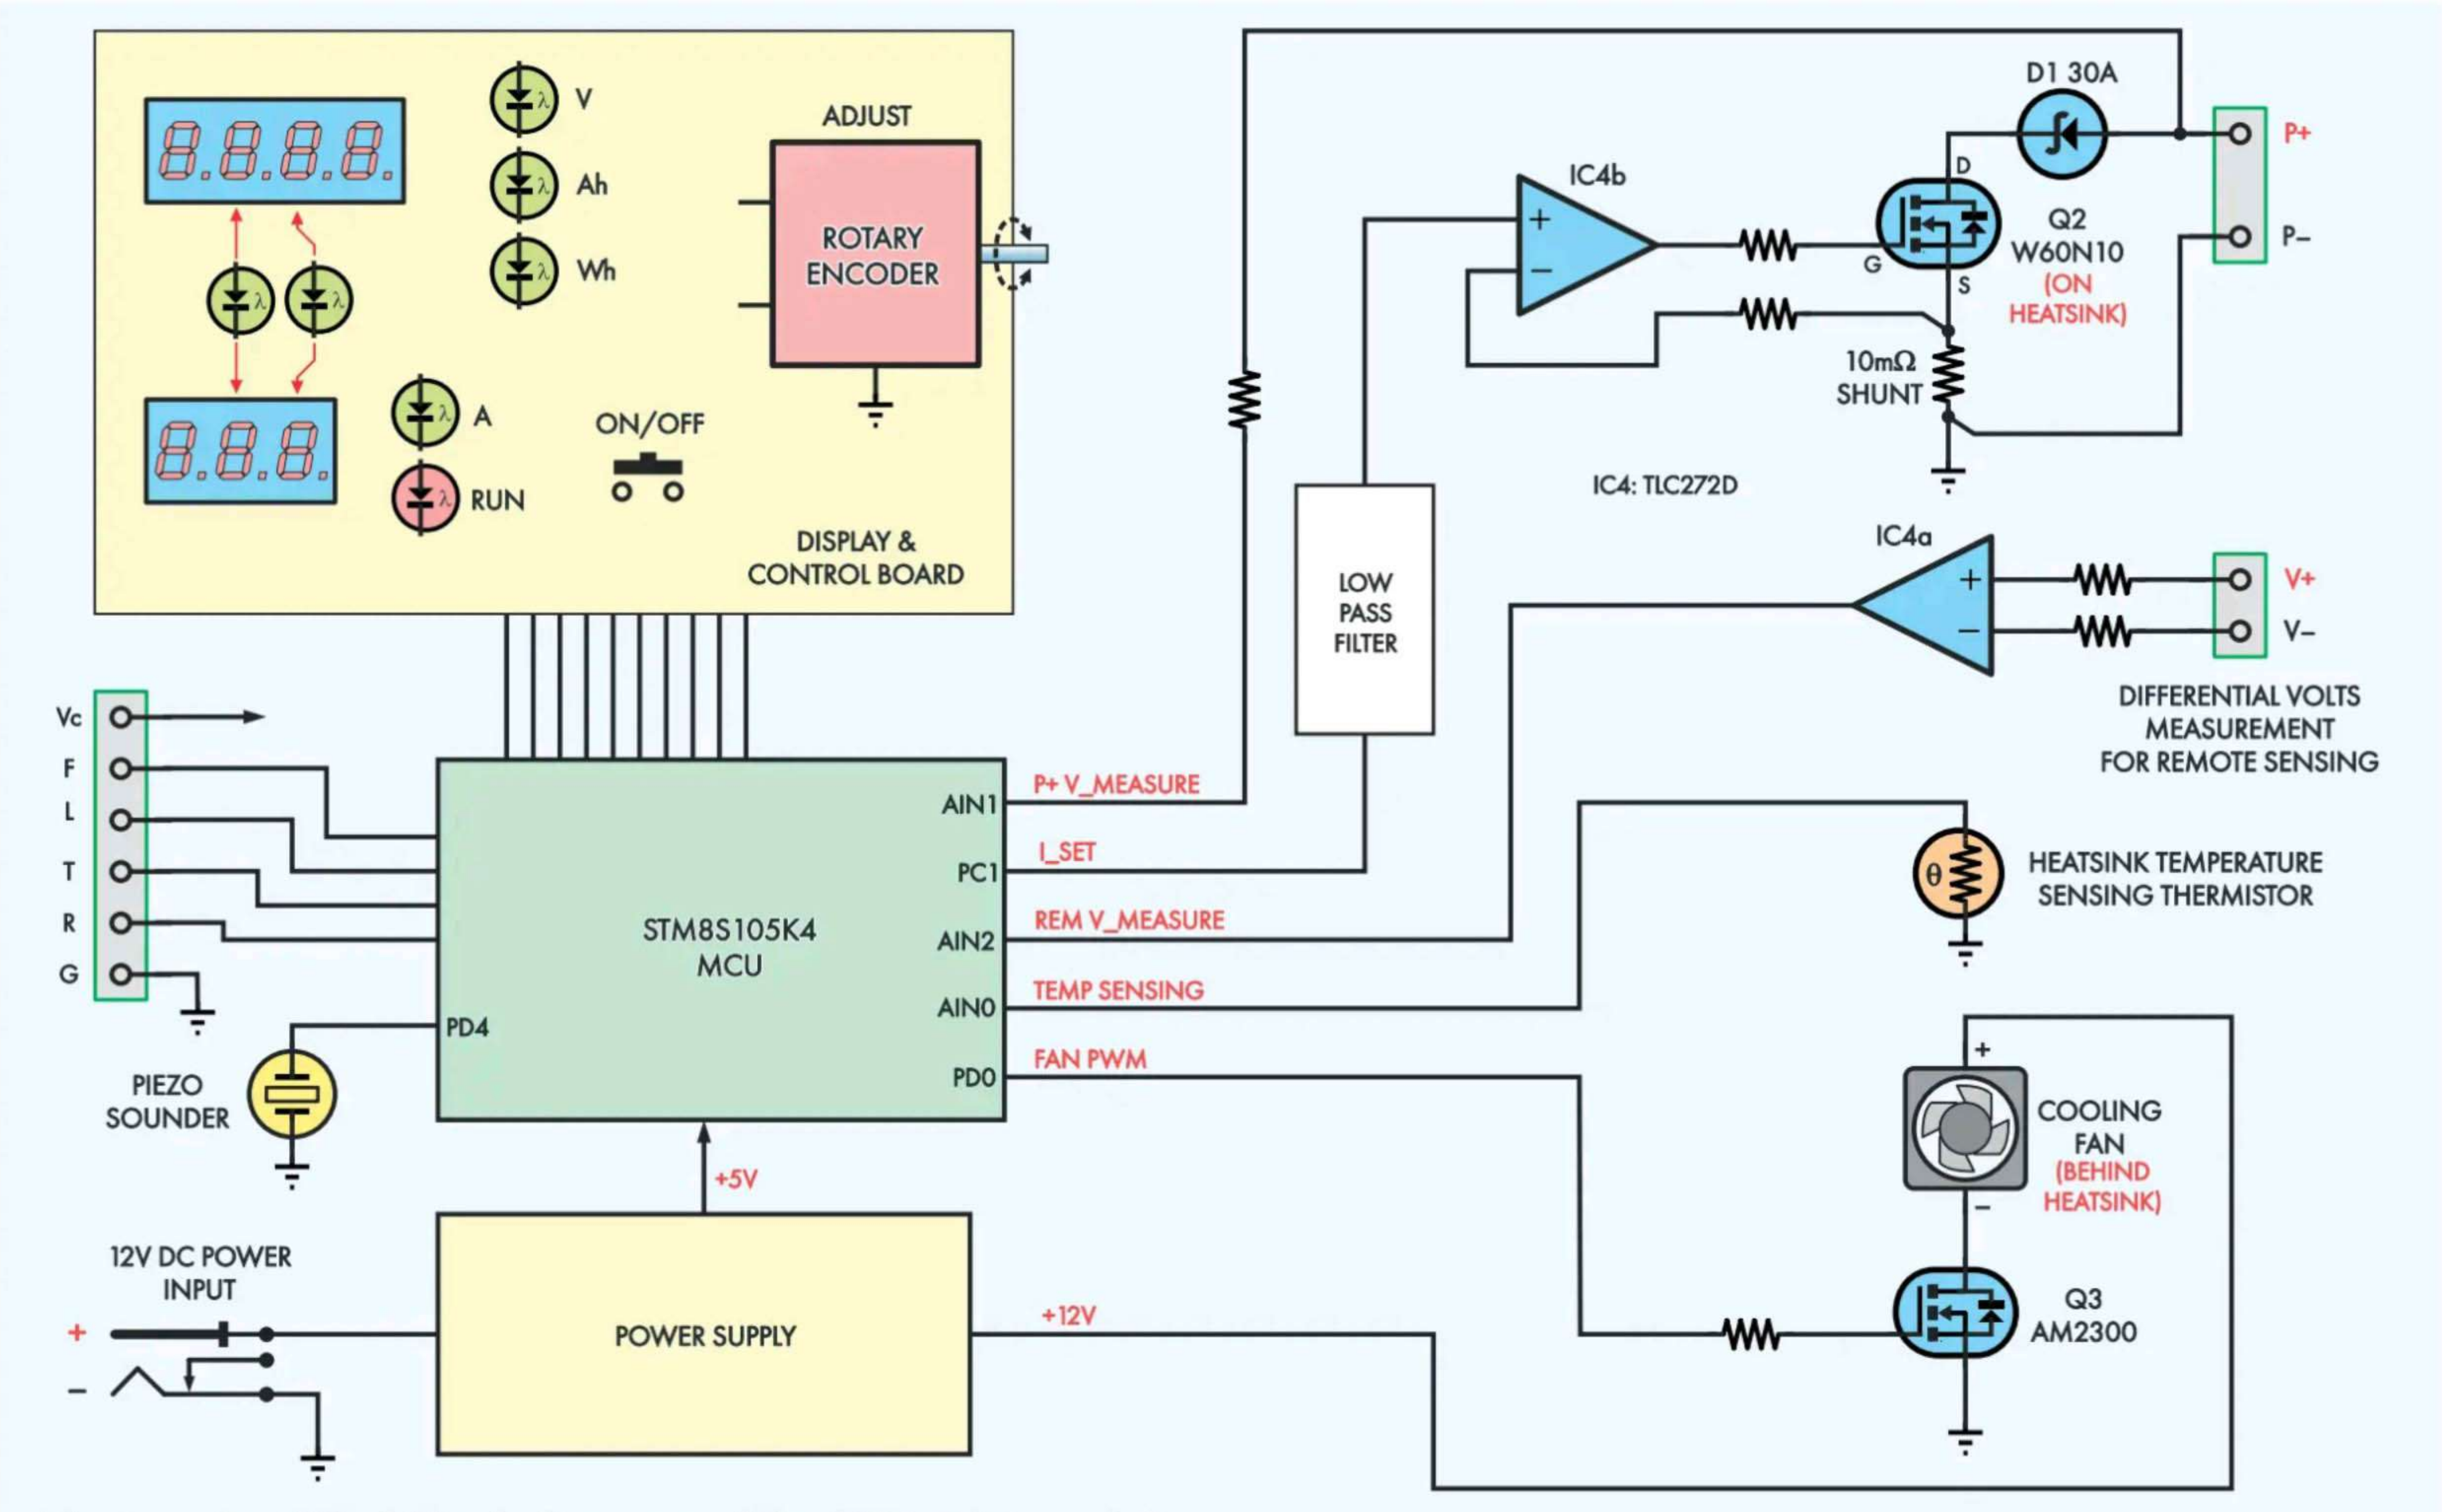
8.8.8.8.
8.8.8.
V
Ah
Wh
A
RUN
ADJUST
ROTARY ENCODER
ON/OFF
DISPLAY & CONTROL BOARD
Vc
F
L
T
R
G
STM8S105K4 MCU
PD4
AIN1
PC1
AIN2
AIN0
PD0
P+ V_MEASURE
I_SET
REM V_MEASURE
TEMP SENSING
FAN PWM
PIEZO SOUNDER
+5V
12V DC POWER INPUT
+
–
POWER SUPPLY
+12V
LOW PASS FILTER
IC4b
IC4: TLC272D
D1 30A
D
G
S
Q2
W60N10
(ON HEATSINK)
10mΩ SHUNT
P+
P–
IC4a
V+
V–
DIFFERENTIAL VOLTS MEASUREMENT FOR REMOTE SENSING
HEATSINK TEMPERATURE SENSING THERMISTOR
COOLING FAN (BEHIND HEATSINK)
Q3
AM2300

Fig.1: a simplified block diagram of the ZPB30A1 module.

powered by a 12V DC supply that must be separate from the measurement current source. The supply voltage must be between 11V and 13V, delivering at least 500mA. Power is applied to the module via a standard barrel-type DC connector (5.5mm outer diameter, 2.1mm inner diameter) on the left side of the main PCB.

When power is applied, the module powers up in whichever of its two operating modes was last used. This mode is displayed in the upper four-digit 7-segment LED display as either “Fun1” for programmable load mode or “Fun2” for battery capacity mode.

The module defaults initially to the Fun1 mode. If you want to switch it to the other mode, you need to switch off the power, wait a few seconds and then hold down the on/off pushbutton while re-applying power. The module then allows you to switch modes using the rotary encoder, after which you press the on/off button again to lock the module into that mode.

The module’s testing inputs are on the right-hand side of the main PCB. The small two-way screw terminal block is the main test input connector, with its inputs labelled “P+” and “P–”. The smaller two-pin socket is used for the optional remote voltage sensing, to avoid errors due to voltage drops in the connecting wires. Its pins are labelled “V+” and “V–”.

How it works

Fig.1 shows a simplified block diagram of the ZPB30A1. I would have liked to show a complete schematic, but all I could find online was a partial circuit (at www.voltlog.com) that had been ‘reverse engineered’ and didn’t cover everything on the main PCB, let alone any of the circuitry on the display/control PCB. Still, all the most important details are shown in Fig.1.

The ‘brains’ of the device is an STM8S105K4 microcontroller unit (MCU), shown at lower left. This responds to the controls on the display and control board shown at upper left and shows the parameter values and testing status on the same board.

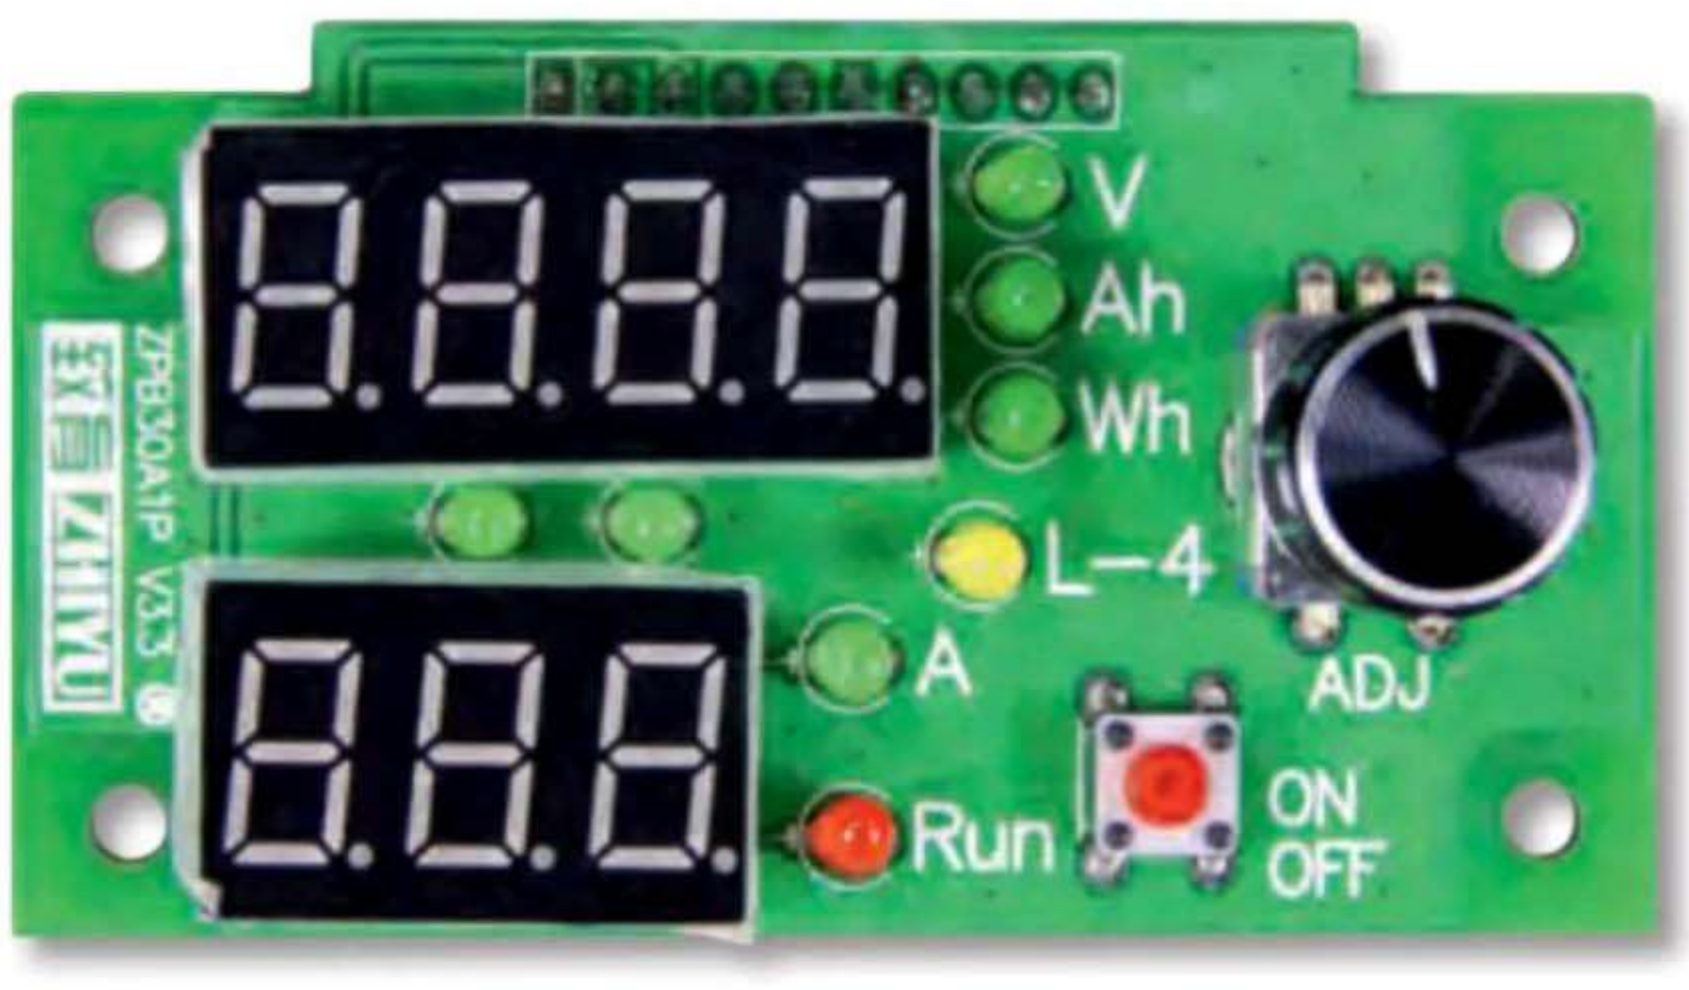
8.8.8.8.
8.8.8.
V
Ah
Wh
L-4
A
Run
ADJ
ON OFF
ZPB30A1P V3.3

The rotary encoder on the display PCB is used for changing modes and adjusting the current & voltage values.

The load current control circuit is shown at upper right. This uses the W60N10 power Mosfet (Q2) to maintain the load current between the test terminals P+ and P–, under the control of the MCU via the I_SET line. The current is monitored using a 10mΩ shunt resistor in Q2’s source connection; op amp IC4b compares the voltage drop across the shunt with the control voltage from the MCU.

You can see a simplified version of the remote differential voltage sensing input below the current control circuit, using op amp IC4a. Its output is taken to the AIN2 analog input of the MCU.

The thermistor mounted next to the Mosfet on the heatsink is shown below the remote voltage sensing input in Fig.1. The TEMP SENSING line from the thermistor goes to the MCU’s AIN0 analog input.

Note that the micro doesn’t have a way to monitor the actual load current – there is no connection from the 10mΩ shunt to the micro. The actual load current will equal the set current almost all the time; if the source

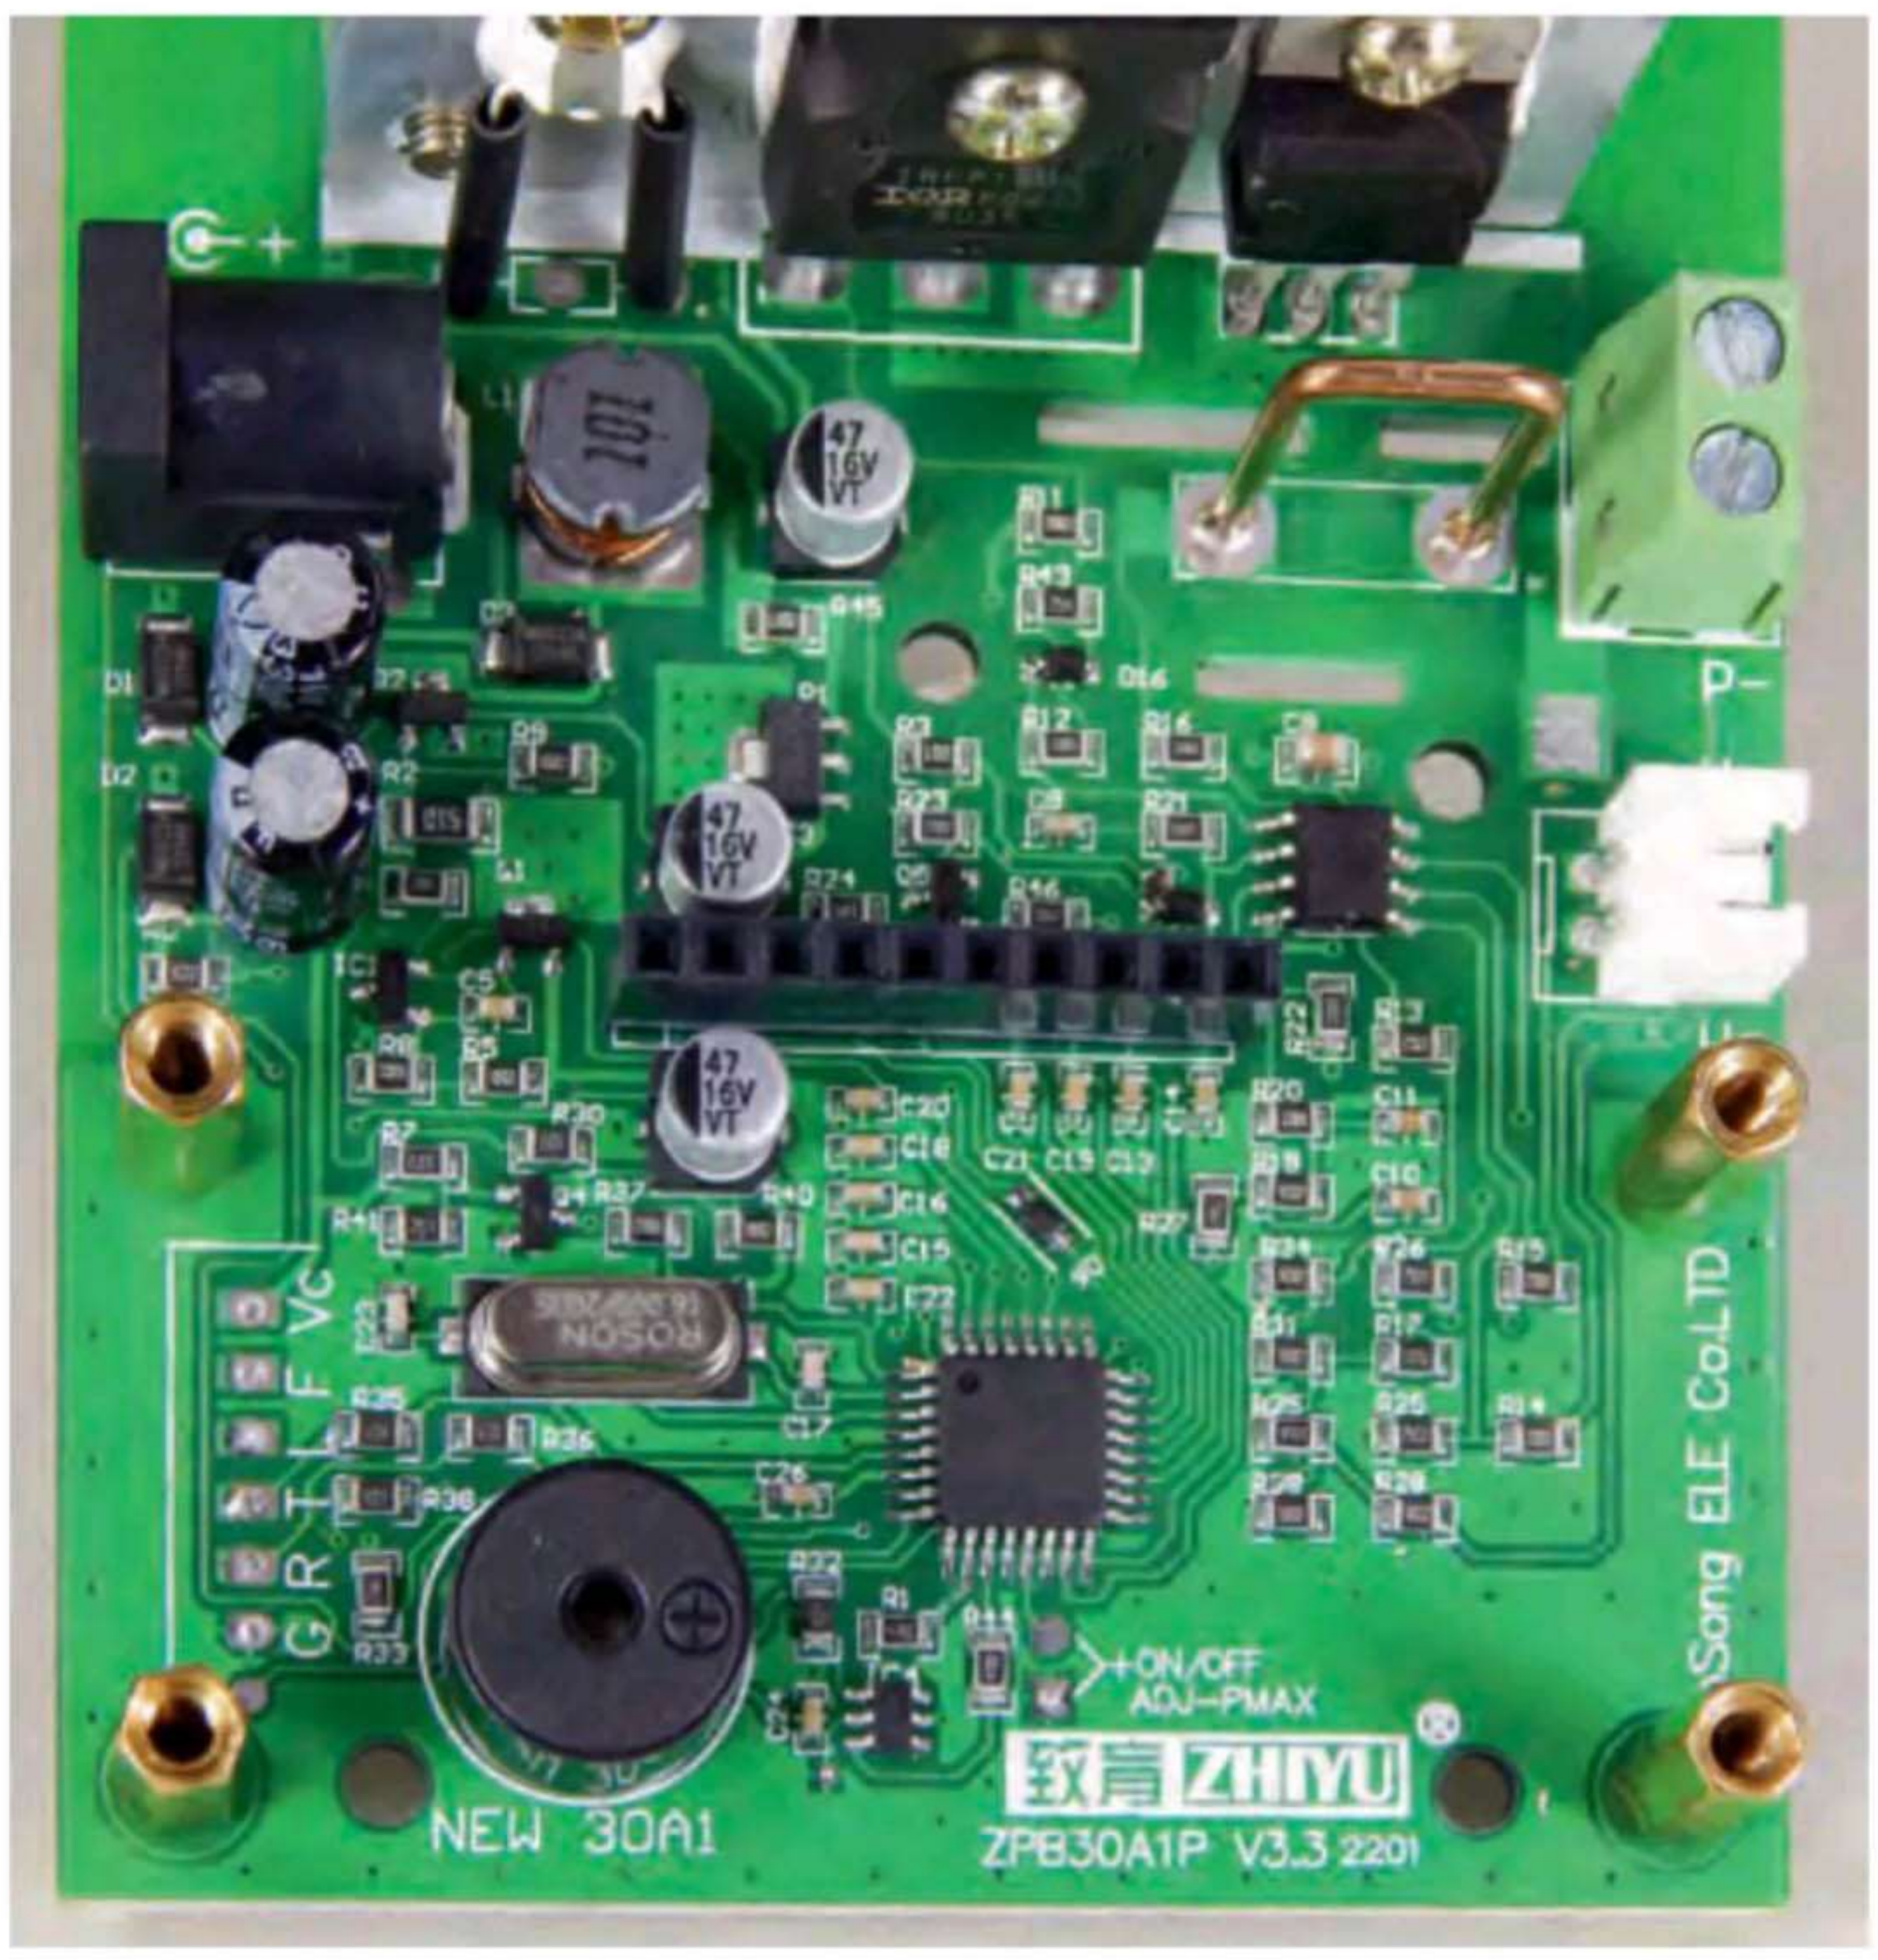

The "main" PCB measures 100 x 69mm and has the heatsink mounted on it. It's also where the majority of the components and power socket are located.

cannot supply enough current to meet the target, the voltage will drop to near-zero, triggering the under-voltage alarm. So it is a safe assumption.

The cooling fan's driver Mosfet, Q3, is controlled by a PWM (pulse-width modulated) signal from the PD0 digital output pin of the MCU. This allows the MCU to turn on the fan as soon as the thermistor reports that the heatsink temperature has risen significantly, and to increase the fan's speed as necessary to keep the temperature under control.

If the temperature keeps rising beyond a safe level, the MCU turns off the load current and stops the test.

The piezo sounder is driven by the MCU's PD4 digital output pin. This allows the MCU to attract your attention whenever it needs to do so; for example, when a test comes to an end, or it detects an error condition.

The module I received did not have the 6-pin header fitted (shown above the piezo sounder); there was just a set of pads and holes labelled G, R, T, L, F and Vc.

While there was no mention of these in the sketchy data provided on the Banggood website, when I searched the internet, I found a couple of suggestions that the G, R and T pins could be used for serial communication with the MCU, at a rate of 115,200 baud and with the standard 8N1 protocol.

The information I found said that the module only transmitted serial data in programmable current mode (Fun1), containing three-byte messages with the first two bytes representing the voltage while the third byte indicated testing status (1 = OK, 0 = undervoltage alarm).

Trying it out

The information on using the module provided on just about all of the supplier websites is very vague and quite hard to follow. As a result, you are largely 'on your own' when it comes to using it. It's a matter of trial and error, not made easy by the multiple functions of the module's controls and LED displays.

That is a pity, since it performs surprisingly well when you manage to get it doing what you want.

The first thing I did was ensure that my module was set to constant-current load mode (Fun1). Then I used the rotary encoder to set the load current for the test. This can be any value between 0.1A (100mA) and 9.99A, in steps of 0.01A (10mA).

After this, I connected the module's P+ and P– terminals to a 0-30V/5A programmable power supply, with one high-resolution bench DMM (digital multimeter) monitoring the current and another monitoring the actual voltage at the P+ and P– terminals. Then I pressed the module's on/off button to begin testing.

I set the power supply to a range of voltage levels (3.30V, 5.00V, 9.00V, 12.00V, 15.00V, 20.00V, 25.00V and 30.00V), and at each voltage level, I set the module to draw a series of current levels. At every current level, I used the bench DMM to set the voltage to precisely the desired level and used the other DMM to check the exact current.

The results of these tests are shown in Fig.2. As you can see, in each case, the applied voltage remained constant over a wide range of current levels. That remained true to a point where either the module stopped the test due

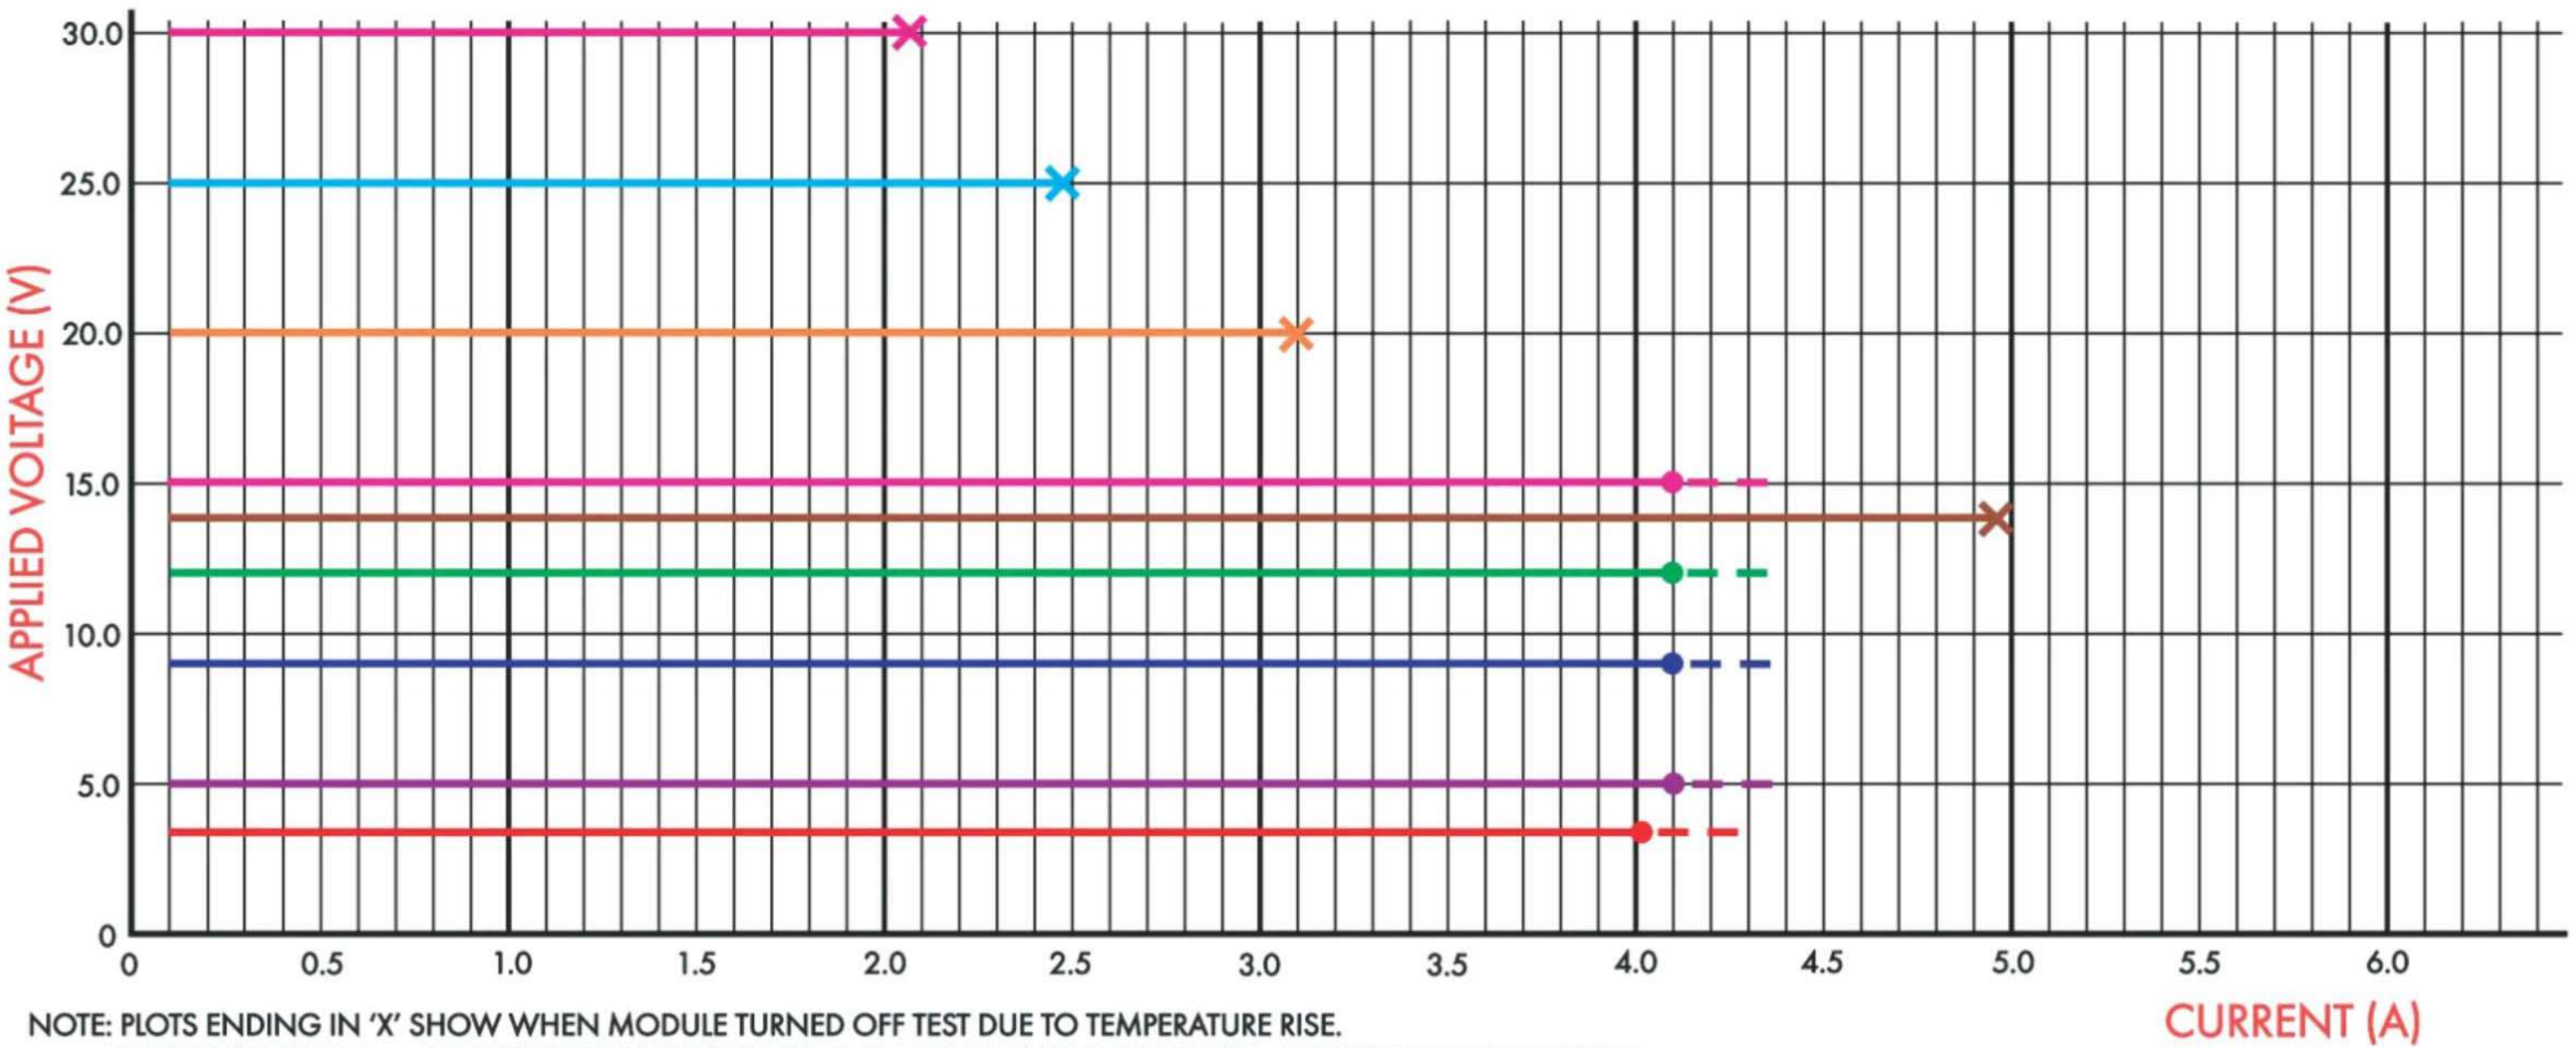

NOTE: PLOTS ENDING IN 'X' SHOW WHEN MODULE TURNED OFF TEST DUE TO TEMPERATURE RISE.
PLOTS ENDING IN ' –•– ' SHOW TESTING NEEDED TO END DUE TO LIMITED CAPABILITY OF TESTING POWER SUPPLY.

Fig.2: I tested the load on constant-current load mode (Fun1) at a range of different voltage levels. This figure shows the current drawn by the module at those voltages.

to the temperature rising above the limit (plots ending in an "X"), or my programmable power supply could provide no more current at that voltage (plots ending in a dot).

Just to make sure, I undertook one further test using a different power supply capable of supplying 13.8V at up to 12A. This resulted in the brown plot in Fig.2. This showed that the module could maintain a current just below 5A at this voltage, corresponding to around 68W dissipation – not bad considering that it is rated to handle a maximum of 60W.

During these tests, I monitored the difference between the module's voltage and current readings and those of the two reference DMMs, to get an idea of the module's measurement accuracy. Its current readings turned out to be less than 0.3% low for currents of 2.0A and above, rising to 1.0% low at 0.5A and 4% down at the lowest current level of 100mA.

These figures compare pretty well with the module's rated accuracy of ±(0.7% + 0.01A).

The voltage readings turned out to be less than 0.4% low over the entire range, which is significantly better than the rated accuracy of ±(1% + 0.02V).

So the ZPB30A1 module performs well in programmable current load (Fun1) mode. I moved on to checking out its battery capacity/Fun2 mode.

Battery capacity testing

I fully charged an 18650 Li-ion cell, then set up the ZPB30A1 module in Fun2 mode with a discharge current of 1.0A and a minimum voltage of 3.00V. After connecting the Li-ion cell to the P+ and P– inputs, I pressed the module's on/off button to begin testing.

Since the module doesn't seem to have any serial output in this mode, I had to record the time and battery voltage the old-fashioned way, using a pen and paper while reading a stopwatch.

The results of this first test are shown in Fig.3 (red plot). As you can see, the cell didn't last all that long at the 1A discharge rate, with its voltage dropping below 3V after only 41 minutes.

The module then displayed its capacity as 0.679Ah, close to my calculated figure of 683mAh (1A × 41 minutes ÷ 60 minutes). So its measurement was only about 0.58% low.

I recharged the same 18650 cell overnight and set the ZPB30A1 to perform a second test at 500mA. I then spent the next few hours recording the battery voltage every five minutes, again in the old-fashioned way.

The results of this second test are shown in the blue plot in Fig.3. It lasted a lot longer this time, with its voltage only reaching just below the test cutoff voltage of 3V after 228 minutes, corresponding to a capacity of 1900mAh. So it's pretty clear that this particular 18650 battery is only capable of delivering its rated capacity at load currents of 500mA or less.

It's also apparent that the ZPB30A1 is well suited to performing the battery capacity testing role, despite a few minor drawbacks.

Summary

The ZPB30A1 module performs both its main functions – a programmable constant current load and battery capacity tester – very well indeed, especially considering its modest price. But it does have a few failings, including the lack of good instructions.

It's also pretty disappointing that its serial communications are so limited. Having an adequately documented serial connector that worked in all modes and provided a complete set of information would make it much easier to monitor the load voltage and time for each measurement.

Adding a serial port header and supporting MCU firmware should be straightforward and would make things a lot easier, especially when testing a battery's capacity. Hopefully, the module makers will add this serial port feature to it soon, making it a really handy piece of test gear.

Despite that, given its low cost, I still think it is worth getting if you think you will use it. SC

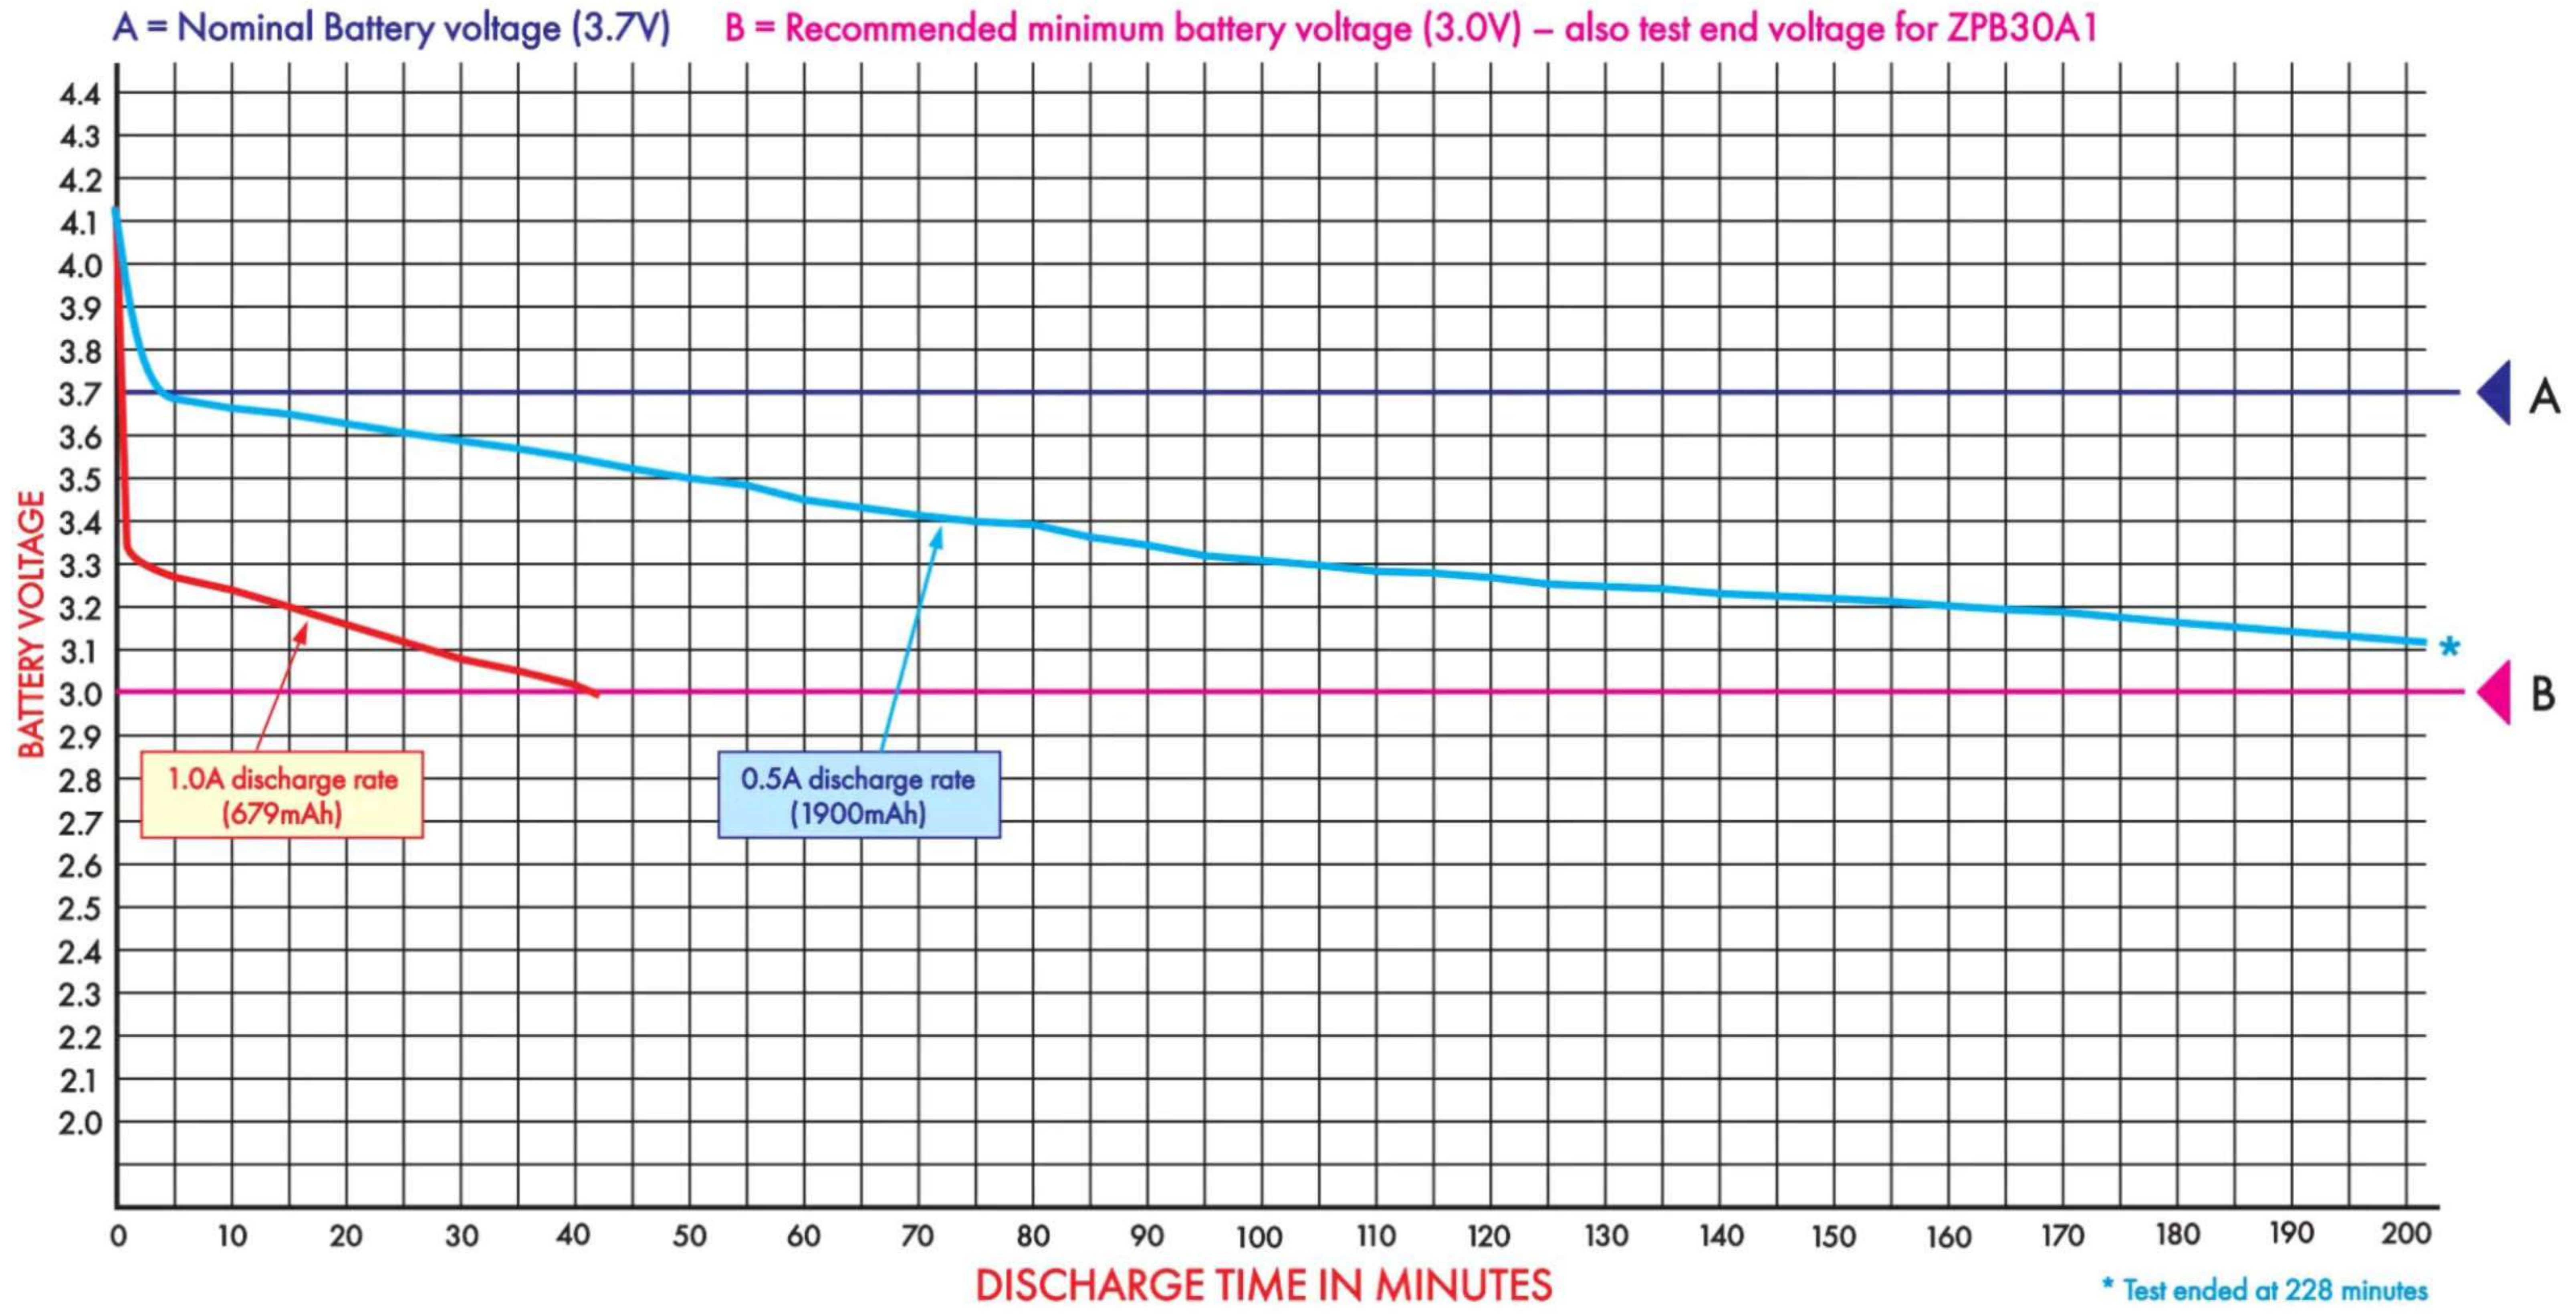

Fig.3: battery capacity testing was performed with a fully-charged 18650 Li-ion cell and the module in Fun2 mode. The test was done with a discharge current of 1A (and later at 0.5A) and battery voltage above 3V.

Part 2 by Tim Blythman

This Multi-Channel Volume Control can handle up to 20 independent channels, allowing you to build your own home theatre or surround system. You can use a touchscreen LCD panel, an IR remote control or an OLED Module with a rotary encoder to control it. This article has all the construction details.

SILICON CHIP
MULTI-CHANNEL VOLUME CONTROL
1.5dB
EEPROM
MUTE
DOWN
UP
MUTE
SETUP

MULTI-CHANNEL VOLUME CONTROL

Our Multi-Channel Volume Control can adjust the levels of up to 20 audio channels by touchscreen, IR remote or a rotary encoder. It's modular, so you can build it with four or eight (or twelve or sixteen) channels if that is all you need.

It's intended to be incorporated as a part of a larger amplifier system, perhaps using several of our Hummingbird Amplifier modules (December 2022). But there is no reason it couldn't be built as a dedicated unit in its own case.

With the principles of operation covered last month, it's time to commence construction. We assume you have already worked out what modules to build and have the parts at hand.

We'll describe the construction of each type of module in turn. You will need one Control and Power Supply Module and at least one Volume Module. If you want a rotary encoder volume control, you must also build the OLED Module.

After that, we'll go over testing the modules and connecting them together into your system. Since all three module types feature surface-mounting parts, check that you have the necessary tools for this sort of work.

A fine-tipped soldering iron, flux paste (and a corresponding cleaning solvent), solder-wicking braid, tweezers and fume extraction are all highly recommended. Some sort of magnifier and a good light source are helpful for those with diminishing eyesight. We're talking from experience here!

Working outside is a good alternative to fume extraction and should also provide sufficient illumination.

Enclosures

If you have decided on your choice of enclosure, you can use the blank PCBs to mark the mounting hole positions. It will be easier to do this now before parts are fitted to the PCBs.

Look at Figs.12 & 13 to get an idea of the cuts that need to be made for the Control and Power Supply Module and OLED Module, respectively.

Control and Power Supply Module

The through-hole parts on this Module are mainly in the power section, while the SMD parts are mostly related to the microcontroller.

We'll start with the SMDs. The three different SOT-23 package parts are the smallest and are all different types, so don't mix them up. Check their orientation against the photos and PCB overlay, Fig.8. SOT-23 parts are small, but their leads are spread out, so you shouldn't get bridges between pins.

REG3 is the MCP1700-3.3 type. Apply flux paste to the three pads and hold it roughly in place with the tweezers. Tack one pin and adjust its position (melting the solder with the iron if needed) until it is square with the pads and flat against the PCB. Then apply solder to each of the other pins in turn. Refresh the first pin if necessary.

Use the same technique to solder the two Mosfets, Q1 and Q2. Q1 is the P-channel part, while Q2 is the N-channel 2N7002 part.

Follow with IC9, the 20-pin PIC16F18146 microcontroller. Check its orientation and ensure that pin 1 is located near the dot on the PCB near where the capacitor will be fitted later.

Like the earlier parts, apply flux and rest the IC in place. It is larger, so it might not need to be held down with tweezers. Tack one lead and adjust the IC location until it is centred on its pads and flat against the PCB.

Next, carefully solder each IC pin to its pad on the PCB. If you do form a solder bridge, leave it for now. Solder the remaining pins to secure the chip in the correct place.

To fix a solder bridge, apply more flux and press the solder braid against the bridge with the iron. When it has drawn up the solder, carefully slide it away from the IC and repeat as necessary.

After using the braid, surface tension should retain enough solder to form a solid joint, as long as the IC is flat against the PCB. If you're not sure, have a close look using a magnifier and refresh the pins with the iron using some more solder and fresh flux.

The remaining surface mounting parts are all M3216/1206 size (3.2×1.6mm) passives and can be fitted using similar techniques. There are two

Multi-channel Volume Control, part 2

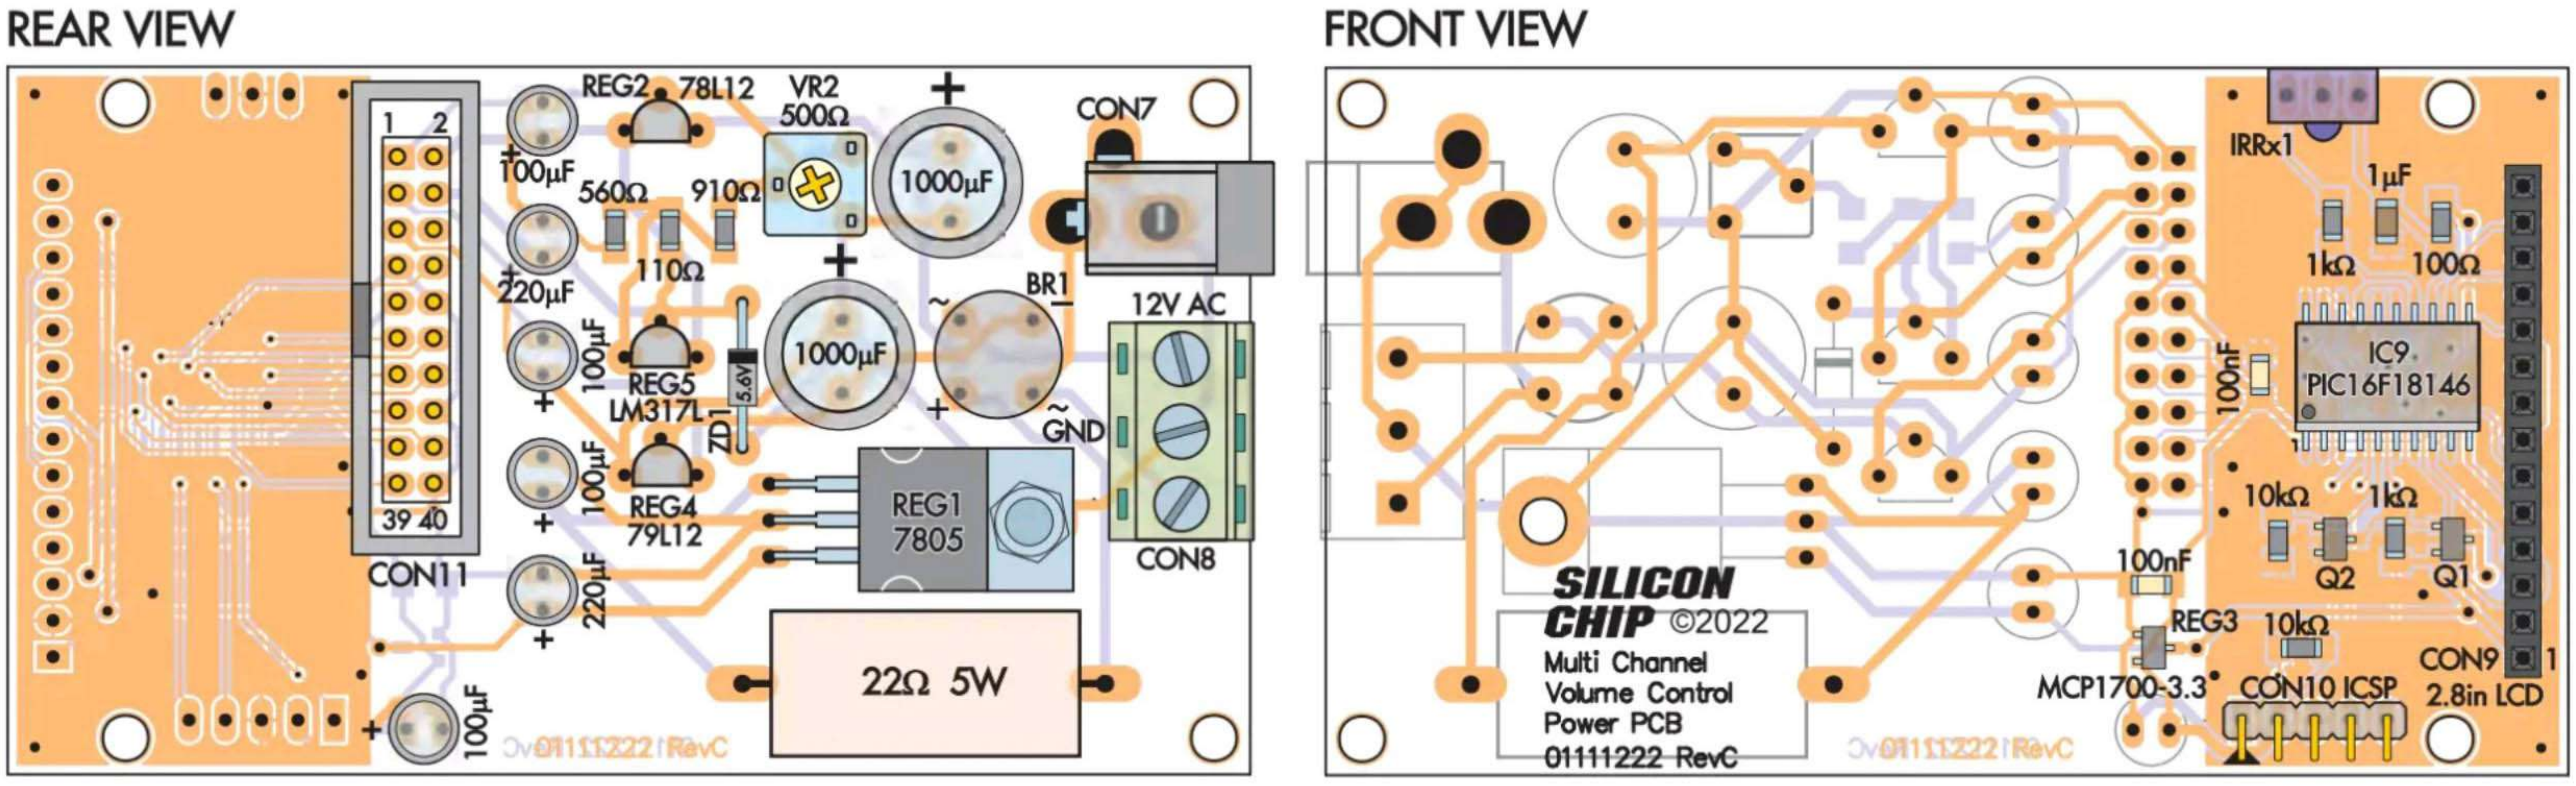
REAR VIEW

FRONT VIEW

REG2 78L12
VR2 500Ω
CON7
100µF
560Ω
910Ω
1000µF
110Ω
220µF
BR1
12V AC
100µF
1000µF
REG5 LM317L
ZD1 5.6V
GND
100µF
REG4 79L12
REG1 7805
CON8
39 40
CON11
220µF
22Ω 5W
100µF
01111222 RevC
IRRx1
1µF
1kΩ
100Ω
IC9 PIC16F18146
100nF
10kΩ
1kΩ
Q2
Q1
SILICON CHIP ©2022
Multi Channel Volume Control Power PCB 01111222 RevC
100nF
REG3
10kΩ
MCP1700-3.3
CON10 ICSP
CON9 2.8in LCD

Fig.8: there is a mix of SMD and through-hole parts, with components on both sides of this Module. Fortunately, none of the SMD parts are too small. Just take care not to mix up the components and watch the orientations of the IC, bridge rectifier, electrolytic capacitors and box header.

100nF capacitors and one 1µF capacitor, which won't be marked, plus some resistors. Five SMD resistors are fitted to the same side of the PCB, plus three on the other side.

To check the resistance codes printed on the parts, refer to Table 1 for the expected markings.

Use a solvent to clean up the excess flux on the PCB. Isopropyl alcohol (isopropanol) is a suitable general-purpose solvent for this. Wipe off as much excess as possible and then allow the remainder to evaporate.

Inspect the PCB with a magnifier to ensure that your soldering looks correct. It will be much easier to make corrections now, before any other components are fitted.

Through-hole parts

The through-hole parts on the Control and Power Supply Module should generally be fitted from shortest to tallest, as that simplifies the process. Refer to the photos and overlay diagrams if needed.

Start with the 5.6V zener diode, ZD1. Bend the leads by 90° and thread through the PCB, ensuring the cathode band matches the PCB silkscreen. Solder the leads and trim so they are neat.

Follow with REG1, the sole TO-220 package regulator. Bend the leads backward by 90° at a point about 7mm from the body. Thread the leads through their PCB pads and affix the regulator with the machine screw, nut and washer. Once you are happy with the location, solder the leads and trim as needed.

Fit bridge rectifier BR1 next, with the + polarity mark on the PCB matching the one on the rectifier. Push it down flat against the PCB before soldering it. Then adjust 500Ω trimpot VR2 so its wiper is at its midpoint and solder it to the PCB.

CON7 and CON8 are next. You don't need to fit both, as only one is needed to supply power, but we used both on our prototype for testing. CON8 is required for a 24V AC centre-tapped supply, while either can be used with a single 12V AC tap. If you have a choice, the 24V AC centre-tapped transformer with CON8 is preferred.

Install the three different regulators in TO-92 cases next, being careful not to mix them up. REG2 is the 78L12, REG4 the 79L12 and REG5 the LM317L. These are also marked on the PCB silkscreen.

Now mount CON11 with the key in the box header facing away from the other components on this side. There might also be a marker on the box header indicating pin 1, which goes near the top of the PCB.

The three different types of electrolytic capacitors are fitted next. There are four 100µF parts and two 220µF parts around CON11. Ensure that the polarities and values are correct before soldering, with the longer, positive leads towards the + markings on the board.

The polarity of the two larger 1000µF capacitors near BR1 are reversed compared to the others.

The last remaining component on this side of the PCB is the 5W resistor. Bend its leads and fit it to the PCB pads. Space the body of the resistor about 5mm clear of the PCB. You can tack one lead and adjust its position (if necessary) before soldering the other lead.

Components on the other side

To help align CON9 for the LCD touchscreen module, fit the four 12mm

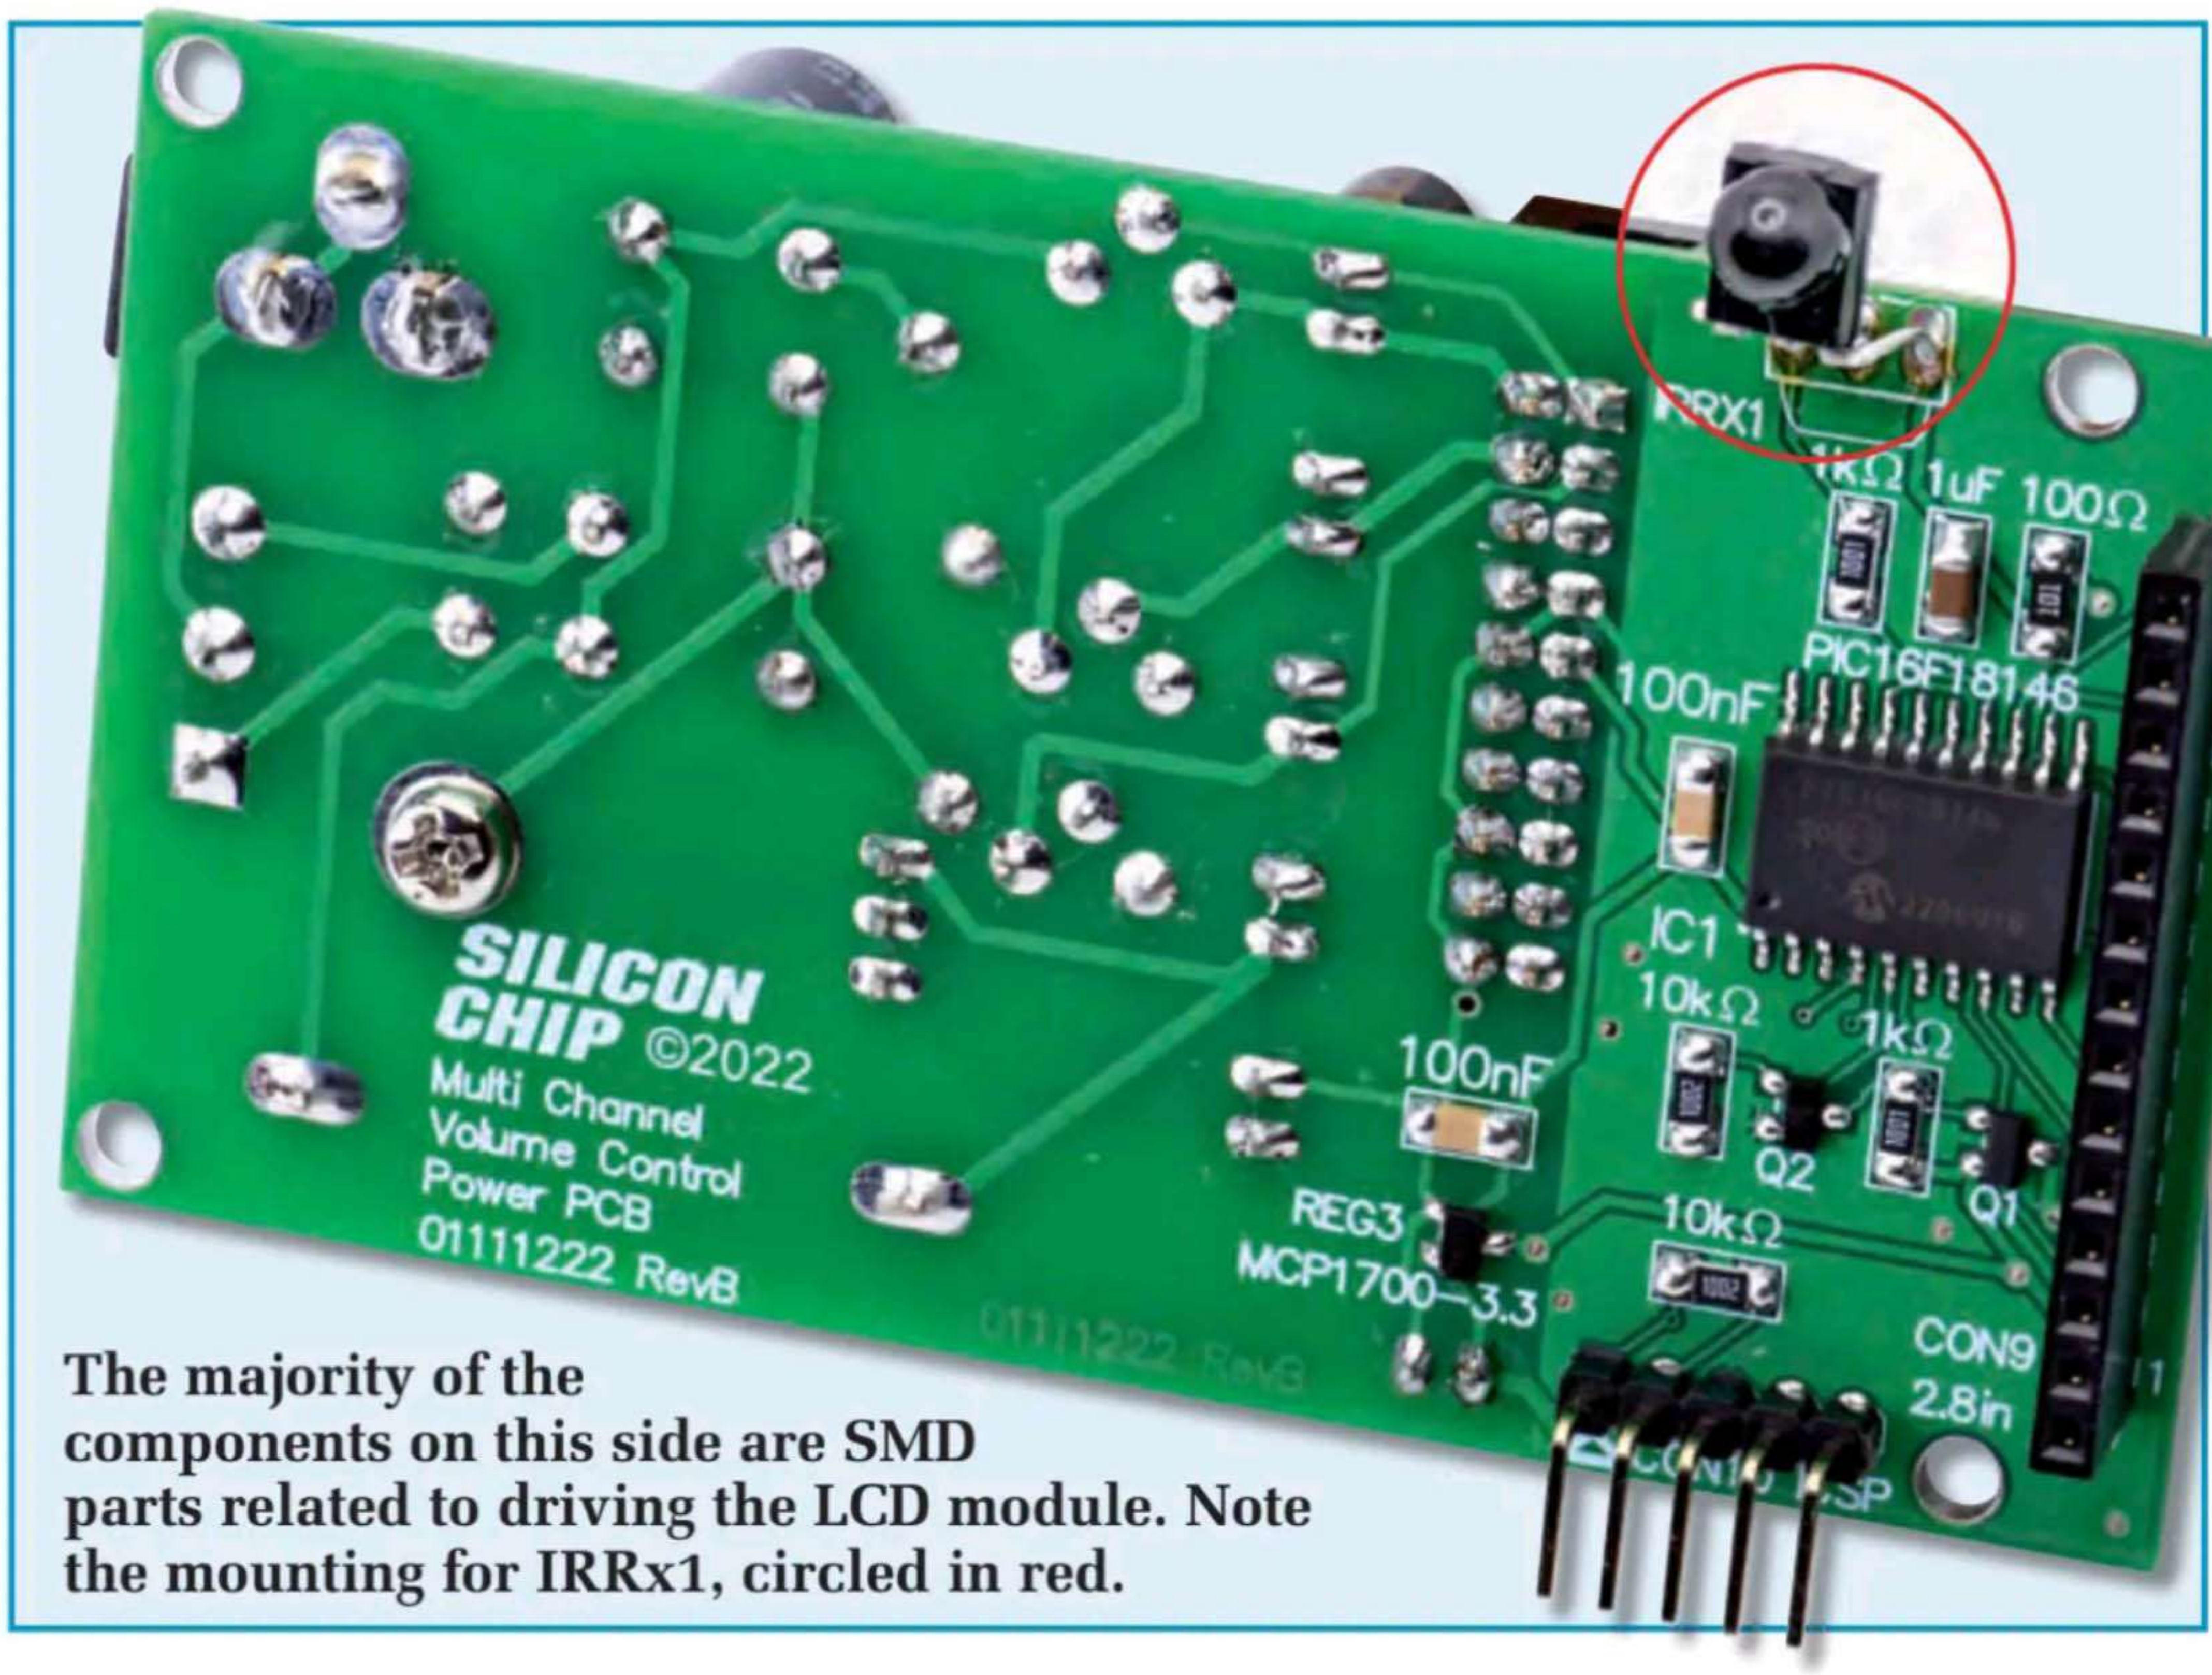
IRX1
1kΩ
1uF
100Ω
100nF
PIC16F18146
IC1
10kΩ
1kΩ
Q2
Q1
SILICON CHIP ©2022
Multi Channel Volume Control Power PCB 01111222 RevB
100nF
REG3
MCP1700-3.3
10kΩ
CON9
2.8in
ICSP

The majority of the components on this side are SMD parts related to driving the LCD module. Note the mounting for IRRx1, circled in red.

M3-tapped spacers to the Module using four M3 screws, with the latter on the same side as the through-hole components.

Rest the 14-way female header (CON9) in place, then slot the LCD module into it, allowing the header to sit at right angles to the PCB. Solder CON9 to the board.

You can test the arrangement for the IR receiver, IRRx1, next. We mounted it so it peeks out just above the top of the LCD (see the photos). There are many ways to mount IRRx1, but we think this method will work in most cases. Regardless of how you do it, just be sure that the correct pins of IRRx1 go to the correct PCB pads.

You should also fit CON10 for in-circuit programming unless you have a pre-programmed microcontroller. We placed it on the top of the PCB, but it could also be fitted to the reverse if necessary.

Programming the micro

If you need to program the microcontroller in-circuit, use a 3.3V supply voltage. Also, detach the LCD module before programming to reduce the load on the programmer's power source.

These newer PICs can only be programmed with a PICkit 4 (or later) or a Snap; with the Snap, you will need to provide power separately. We discussed modifying a Snap to supply power on page 31 of the June 2022 issue (in the PIC Programming Helper project article).

Use the IPE to upload the 0111122B.HEX file (0111122C.HEX is for the OLED Module) and confirm that you get the "Program/Verify complete" message. You won't see anything that indicates that it is working right away.

Testing

Leave the LCD module off when checking the supply rails on the Control and Power Supply Module. It's a good idea to do this with nothing attached, especially as we need to trim the 5.5V rail.

You can connect a current limited DC supply (eg, a bench supply) to the CON8 screw terminals. Connect the negative supply to CON8's centre GND connection with the positive supply to either of the remaining terminals. This will provide power to the positive regulators. Reversing the polarity will power the negative regulator, which we will do later.

Set the current limit to around 100mA and slowly wind up the supply voltage. With 15V applied, we found that our prototype's 12V, 5.5V, 5V and 3.3V rails were correct (within 0.1V), with the Module drawing about 60mA.

You can access the 12V, 5.5V and 3.3V rails at pins 2, 4 and 9 of CON11, respectively. The 5V rail can be sensed at pin 1 of CON9 (where the LCD panel connects). CON8's centre pin or REG1's tab are good places to connect to ground for referencing these readings.

Assuming there is 12V on the 12V rail, adjust VR2 to get a reading of 5.50V, or as close as possible, across ZD1. Don't exceed 5.6V, or ZD1 will start conducting and could get warm.

If you can't trim the 5.5V rail, check the resistor values. Since the other regulators are fixed, there isn't much else that can go wrong apart from the wrong regulator being fitted or the bridge rectifier not being installed correctly.

Reverse the polarity applied to CON8 to check the -12V rail at pin 3 of CON11. In this case, we found that the Module only drew about 30mA.

LCD module backlight

One of the problems we encountered during the design and testing of this project and the earlier Digital Preamp is that the LCD backlight has the heaviest current draw of any component.

In the Digital Preamp, we applied the well-known technique of modulating that draw by applying a PWM signal to the backlight control. For this project, we wanted to tackle this in a better way, as it was apparent that the PWM signal was having a small but noticeable effect on the measured audio quality.

So for this project, we have avoided using PWM control of the LCD backlight. You can see that the power section of the circuit now uses a 5W resistor instead of several 1W resistors, so it is better able to handle the full backlight current.

We still found that the 5W resistor was getting warm, so we had a closer look at what we could do to reduce dissipation. While getting the Module to run cooler is always an advantage, we hoped the lower current draw would lead to less ripple on the main supply capacitors and thus better performance.

The "LCD screen backlight modifications" panel explains how the backlight works on these LCD modules and discusses a minor modification that can be made to reduce its current draw.

This modification is optional, so you can skip it if you like. Reattach and secure the LCD module using the four remaining machine screws. We can now test that the microcontroller is working correctly and can produce a display on the LCD screen.

Table 1 – SMD resistor codes

Value	3-digit code	4-digit code
100kΩ	104	1003
47kΩ	473	4702
22kΩ	223	2202
10kΩ	103	1002
2.2kΩ	222	2201
1kΩ	102	1001
910Ω	911	910R
680Ω	681	680R
560Ω	561	560R
110Ω	111	110R
100Ω	101	100R

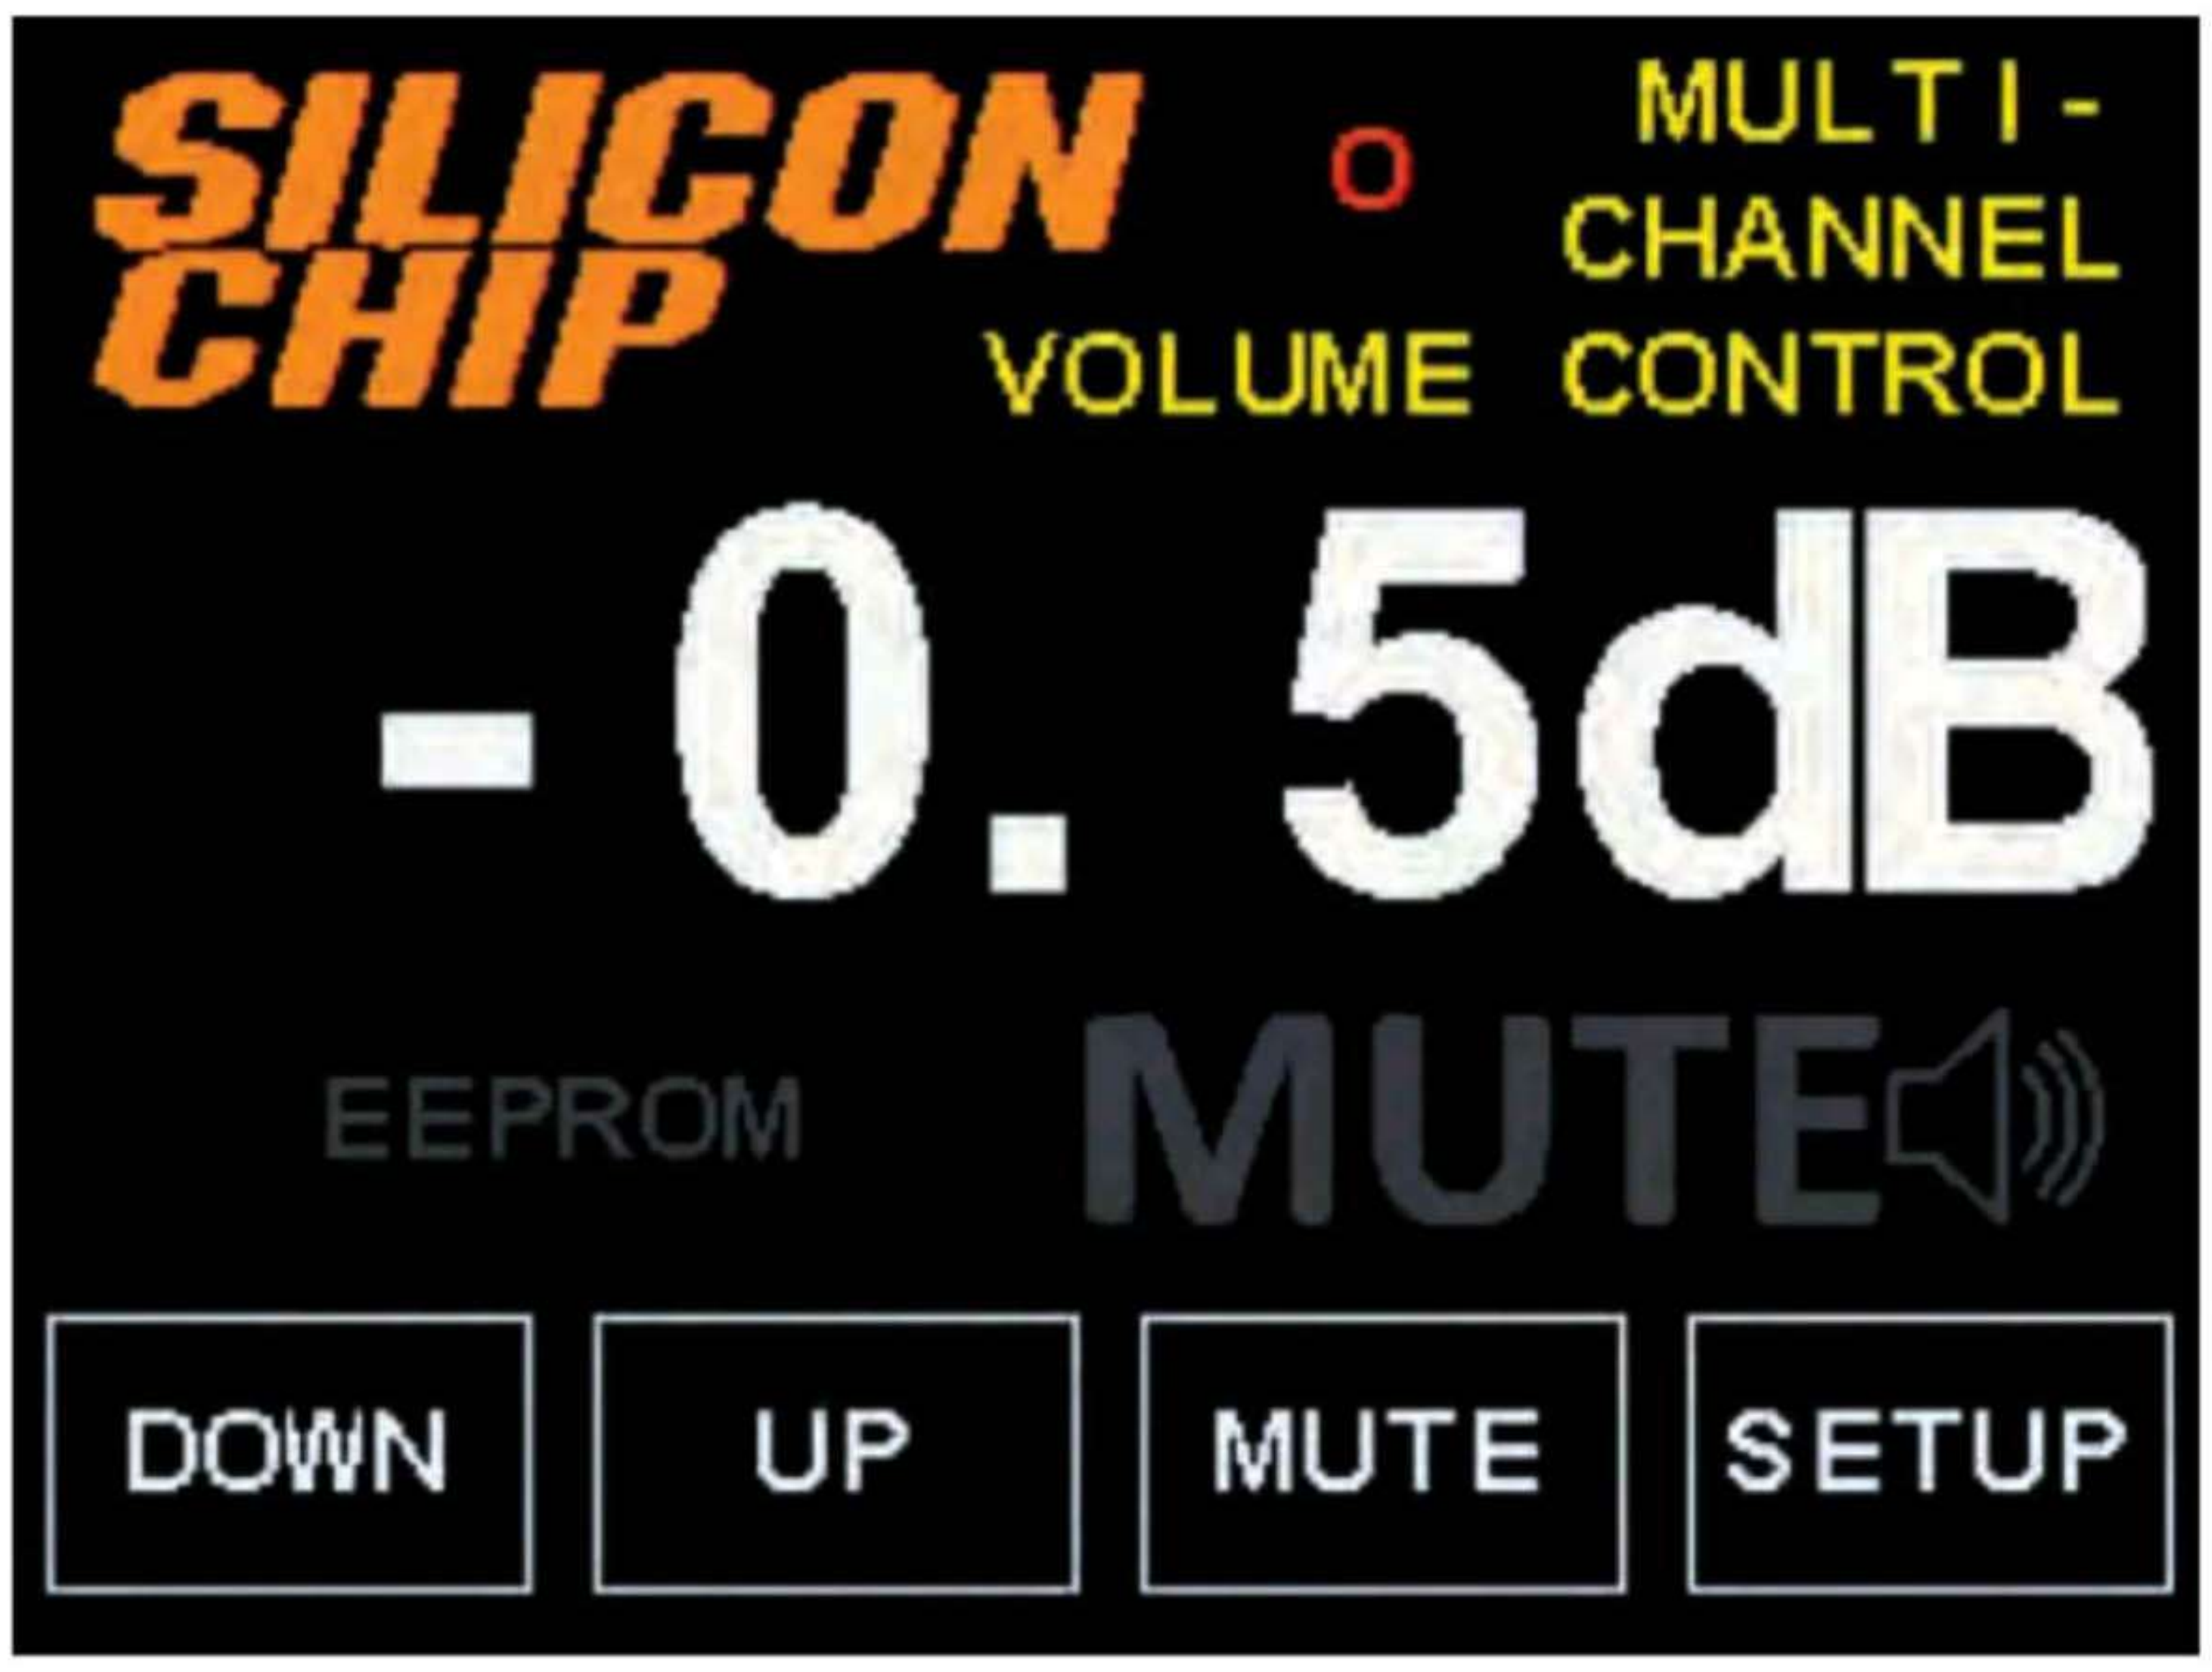

Screen 1: if you see this screen when you power up your Multi-Channel Volume Control, the Control Module is functional. The red circle at upper right is an IR (infrared) telltale that lights up whenever an IR remote control signal is received (whether it is recognised or not).

Multi-channel Volume Control, part 2

Using the connections you used to test the positive regulators (ie negative to GND, positive to either of the AC[~] connections), set the limit to around 300mA and wind up the voltage. You should see something on the LCD with the input at about 8V or higher. If you don't see anything by 15V, there may be a problem.

The actual current draw will depend on the type of LCD backlight and may be different if it has been modified, but it shouldn't be any higher than 300mA.

You should see a screen similar to Screen 1, and the UP/DOWN/MUTE buttons should respond to presses. That's as much as we can test at this point.

Volume Module

The Volume Module can be built without the last op amp stage if you want to save a bit of money and time, and that will also improve the volume control range if you don't need the high maximum gain. All our performance specs are based on the fully populated version; performance will likely be the same or better without those extra op amps.

We'll describe the assembly for all components being fitted. If you wish to leave out the last op amp stage, omit IC3, IC7, their respective 100nF capacitors (one each) and the eight 1kΩ resistors in that area of the PCB.

The two remaining 1kΩ resistors that pad VR1 are at the other end of this PCB. They are still used.

If you omit IC3 and IC7, short out the four PCB jumpers pairs, JP3-JP6. Apply your iron to the pads of the jumper and feed in a generous amount of solder until a bridge forms.

You can use solder wicking braid if you need to remove the bridges.

Fig.9 and the PCB photos show the fully populated version that we will now assemble.

The Volume Module is mostly populated with SMD components, with just a handful of through-hole parts. Start by fitting the eight dual diodes in SOT-23 packages.

Apply flux to the pads and rest each diode in place, noting the orientation from the photos and overlay diagram. Tack one lead, adjust the positioning and then solder the remaining leads. Add some fresh flux and touch the iron to the first lead if you need to refresh that joint.

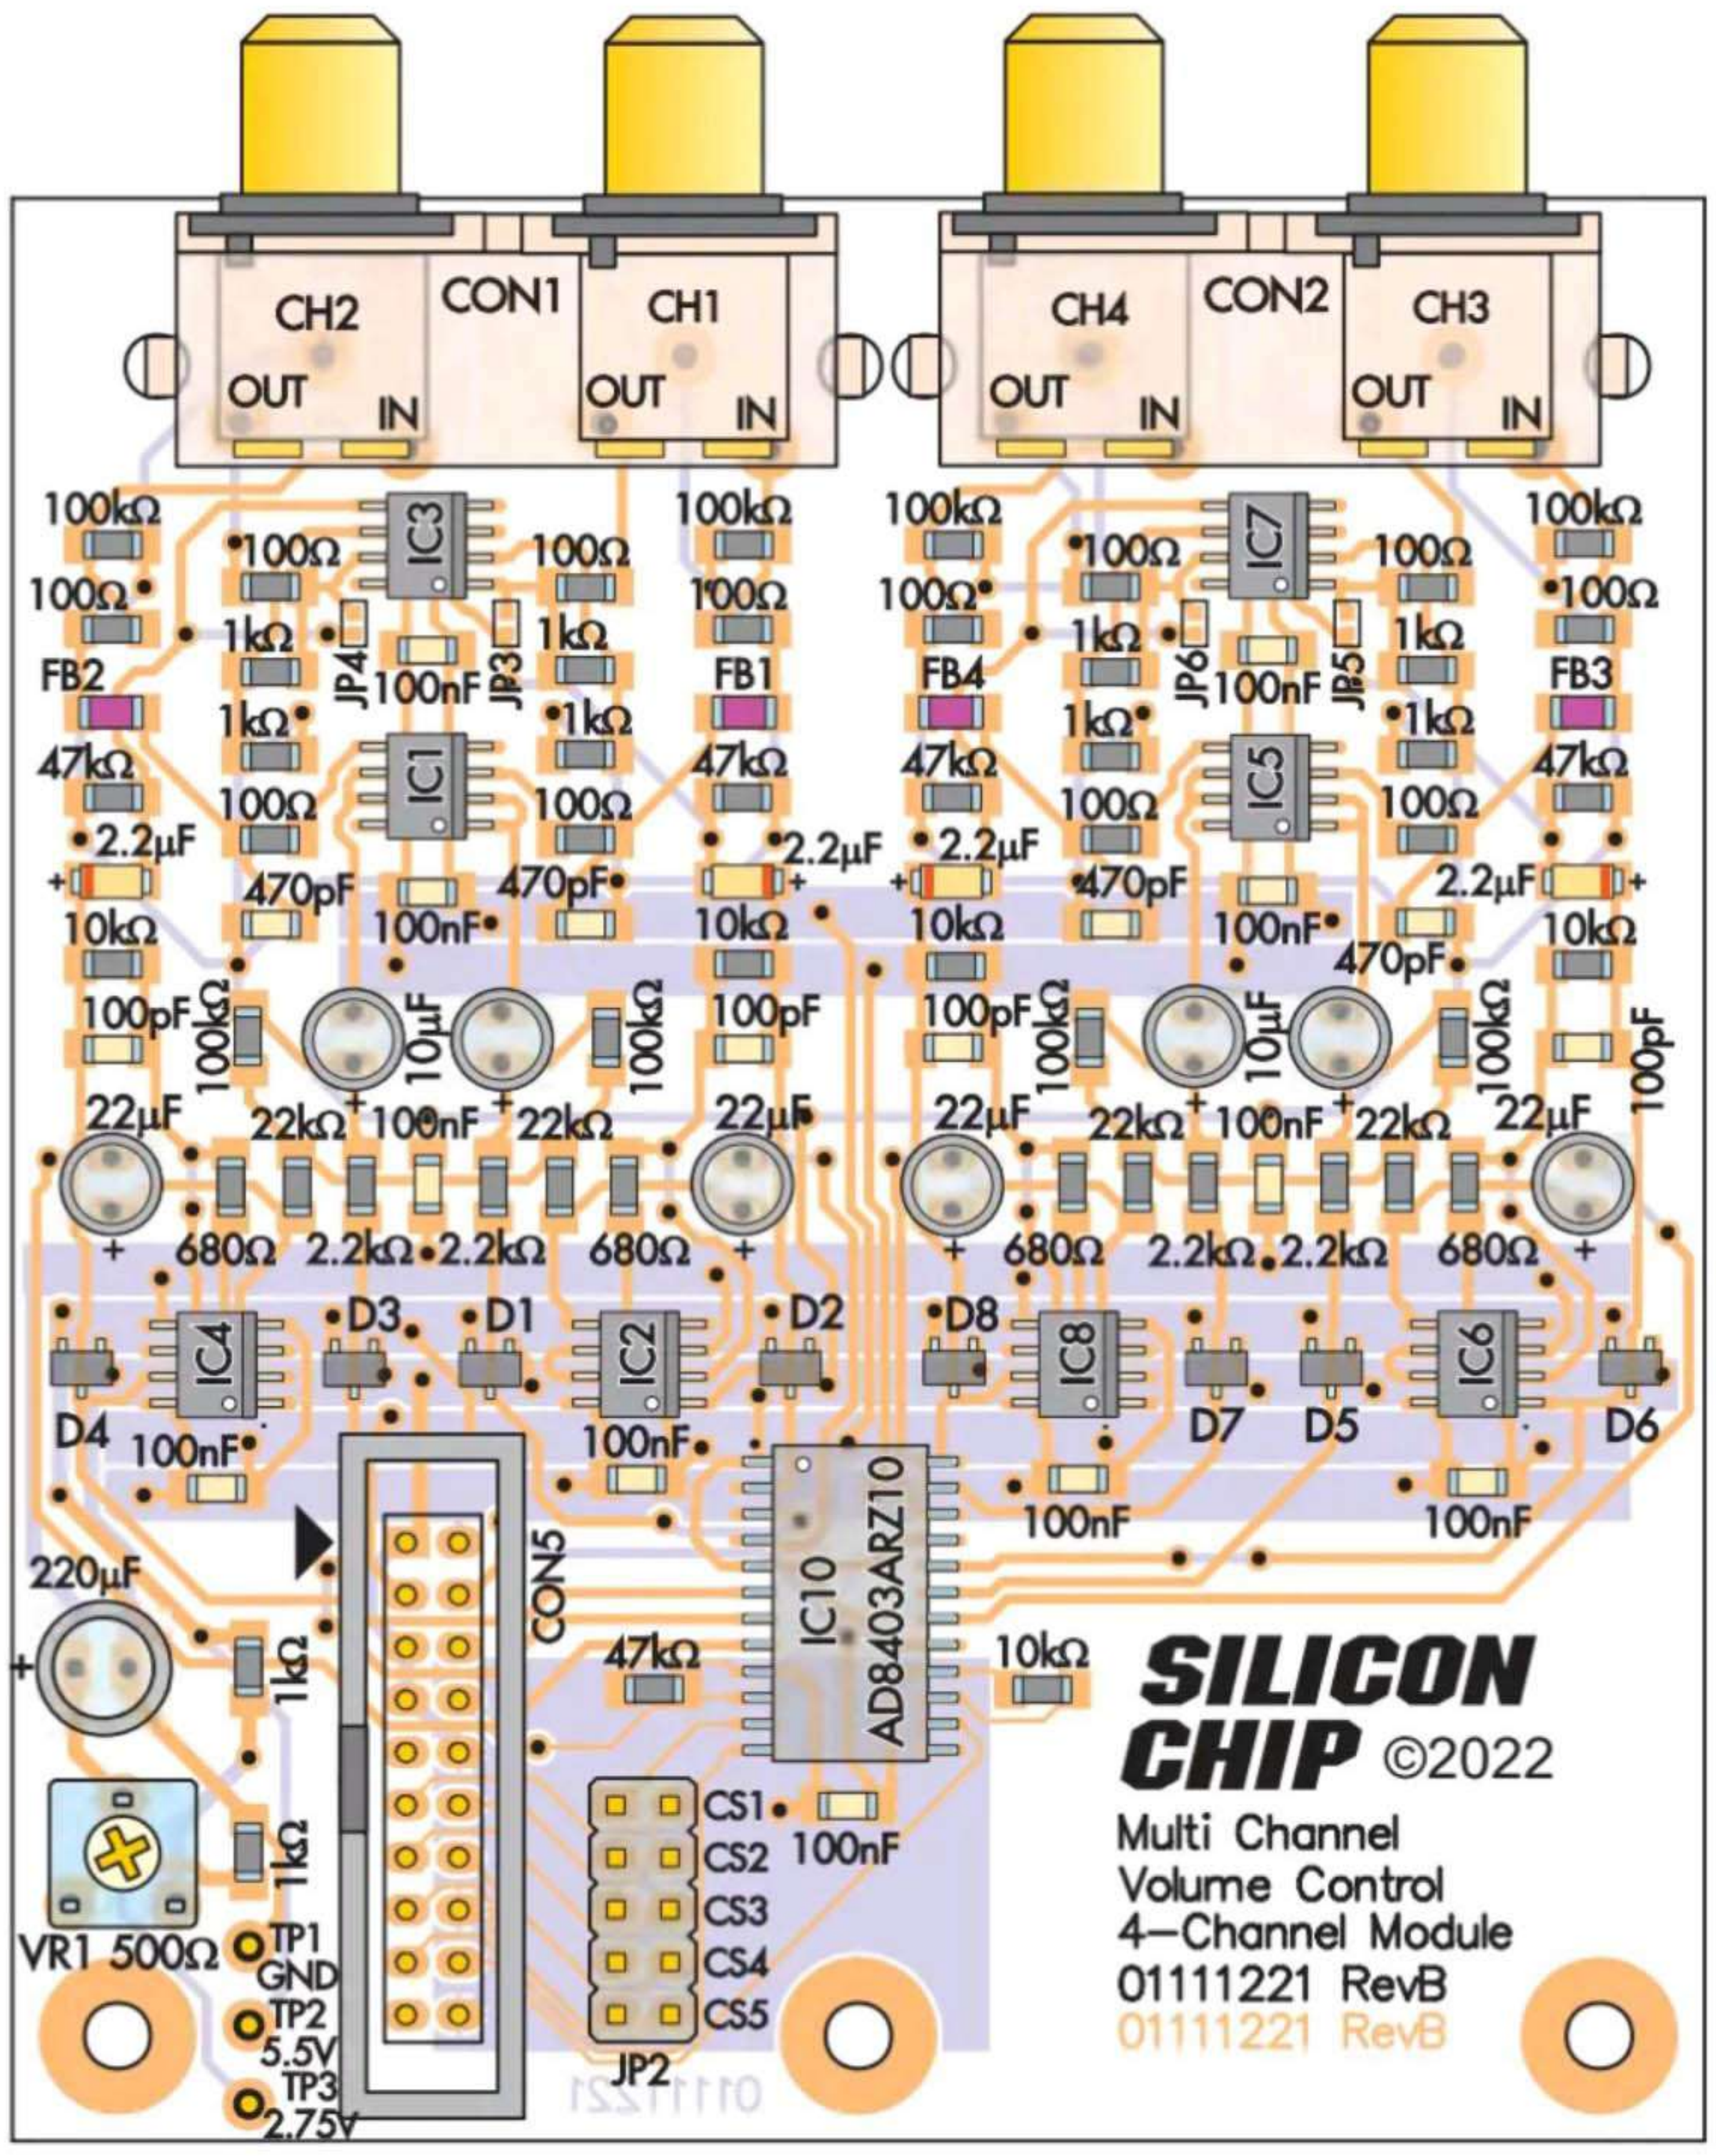

Fig.9: all parts for the Volume Module mount on the top side. Slightly smaller M2012/0805 size passives will fit the same pads. Watch the orientations of the ICs, the electrolytic (including tantalum) capacitors and the box header.

Follow with the eight op amps, IC1-IC8. They all face the same way, with their pin 1 facing towards the bottom of the PCB. Small parts like this may not have a dot printed on their bodies, but may have a bevel along the edge nearest pin 1. This bevel is most easily seen from the end of the chip.

IC10, the 28-pin SOIC part, should be soldered next. Its pin 1 is orientated in the opposite direction from IC1-IC8. If you have any solder bridges on these parts, rectify them using more flux and solder-wicking braid.

The top half of the PCB is marked with horizontal lines and values down the middle, indicating that four identical parts are fitted across. Each part corresponds to one of the four channels, hence the symmetry (Fig.9 shows the values individually for clarity).

The remaining SMD parts on this PCB are two-lead passives. Fit ferrite beads FB1-FB4 next. They are identical and marked as FB on the PCB silkscreen and overlay diagram. The ferrite beads will probably be dark grey, matching the ferrite material they are made from; ceramic capacitors are usually a lighter beige/brown colour.

The four tantalum capacitors are in a row near IC1 and IC5. As these are polarised, observe the polarity markings. It's important to note that, unlike electrolytic can capacitors, rectangular moulded (as well as tag tantalum) electros have a stripe on the positive end, similar to a diode's cathode marking!

Also, our prototype used ceramic capacitors, which look different to the tantalum parts we will supply in kits. You could use high-value SMD ceramics if absolutely necessary, but they are generally inferior for audio signal coupling compared to electrolytic caps.

After that, install the 11 100nF capacitors, the four 470pF capacitors and the 100pF capacitors.

There are many SMD resistors of different values; naturally, they should not be mixed up. Fortunately, their values will be marked, so you can check them as you go (you might need a magnifier) – see Table 1. If you're unsure of reading the codes, carefully use a multimeter to measure their resistances.

You could even measure them using our Advanced SMD Test Tweezers from the February and March 2024 issues.

Cleaning and checking

Now use your preferred flux cleaning solvent to remove any excess flux from the PCB and allow it to dry. It's a good time to inspect the assembly and check that all the components look to be soldered correctly in the right spots before fitting the remaining components.

For JP2, you might like to use a simple wire link if you know what your configuration will be. If so, populate the first board with a link across CS1,

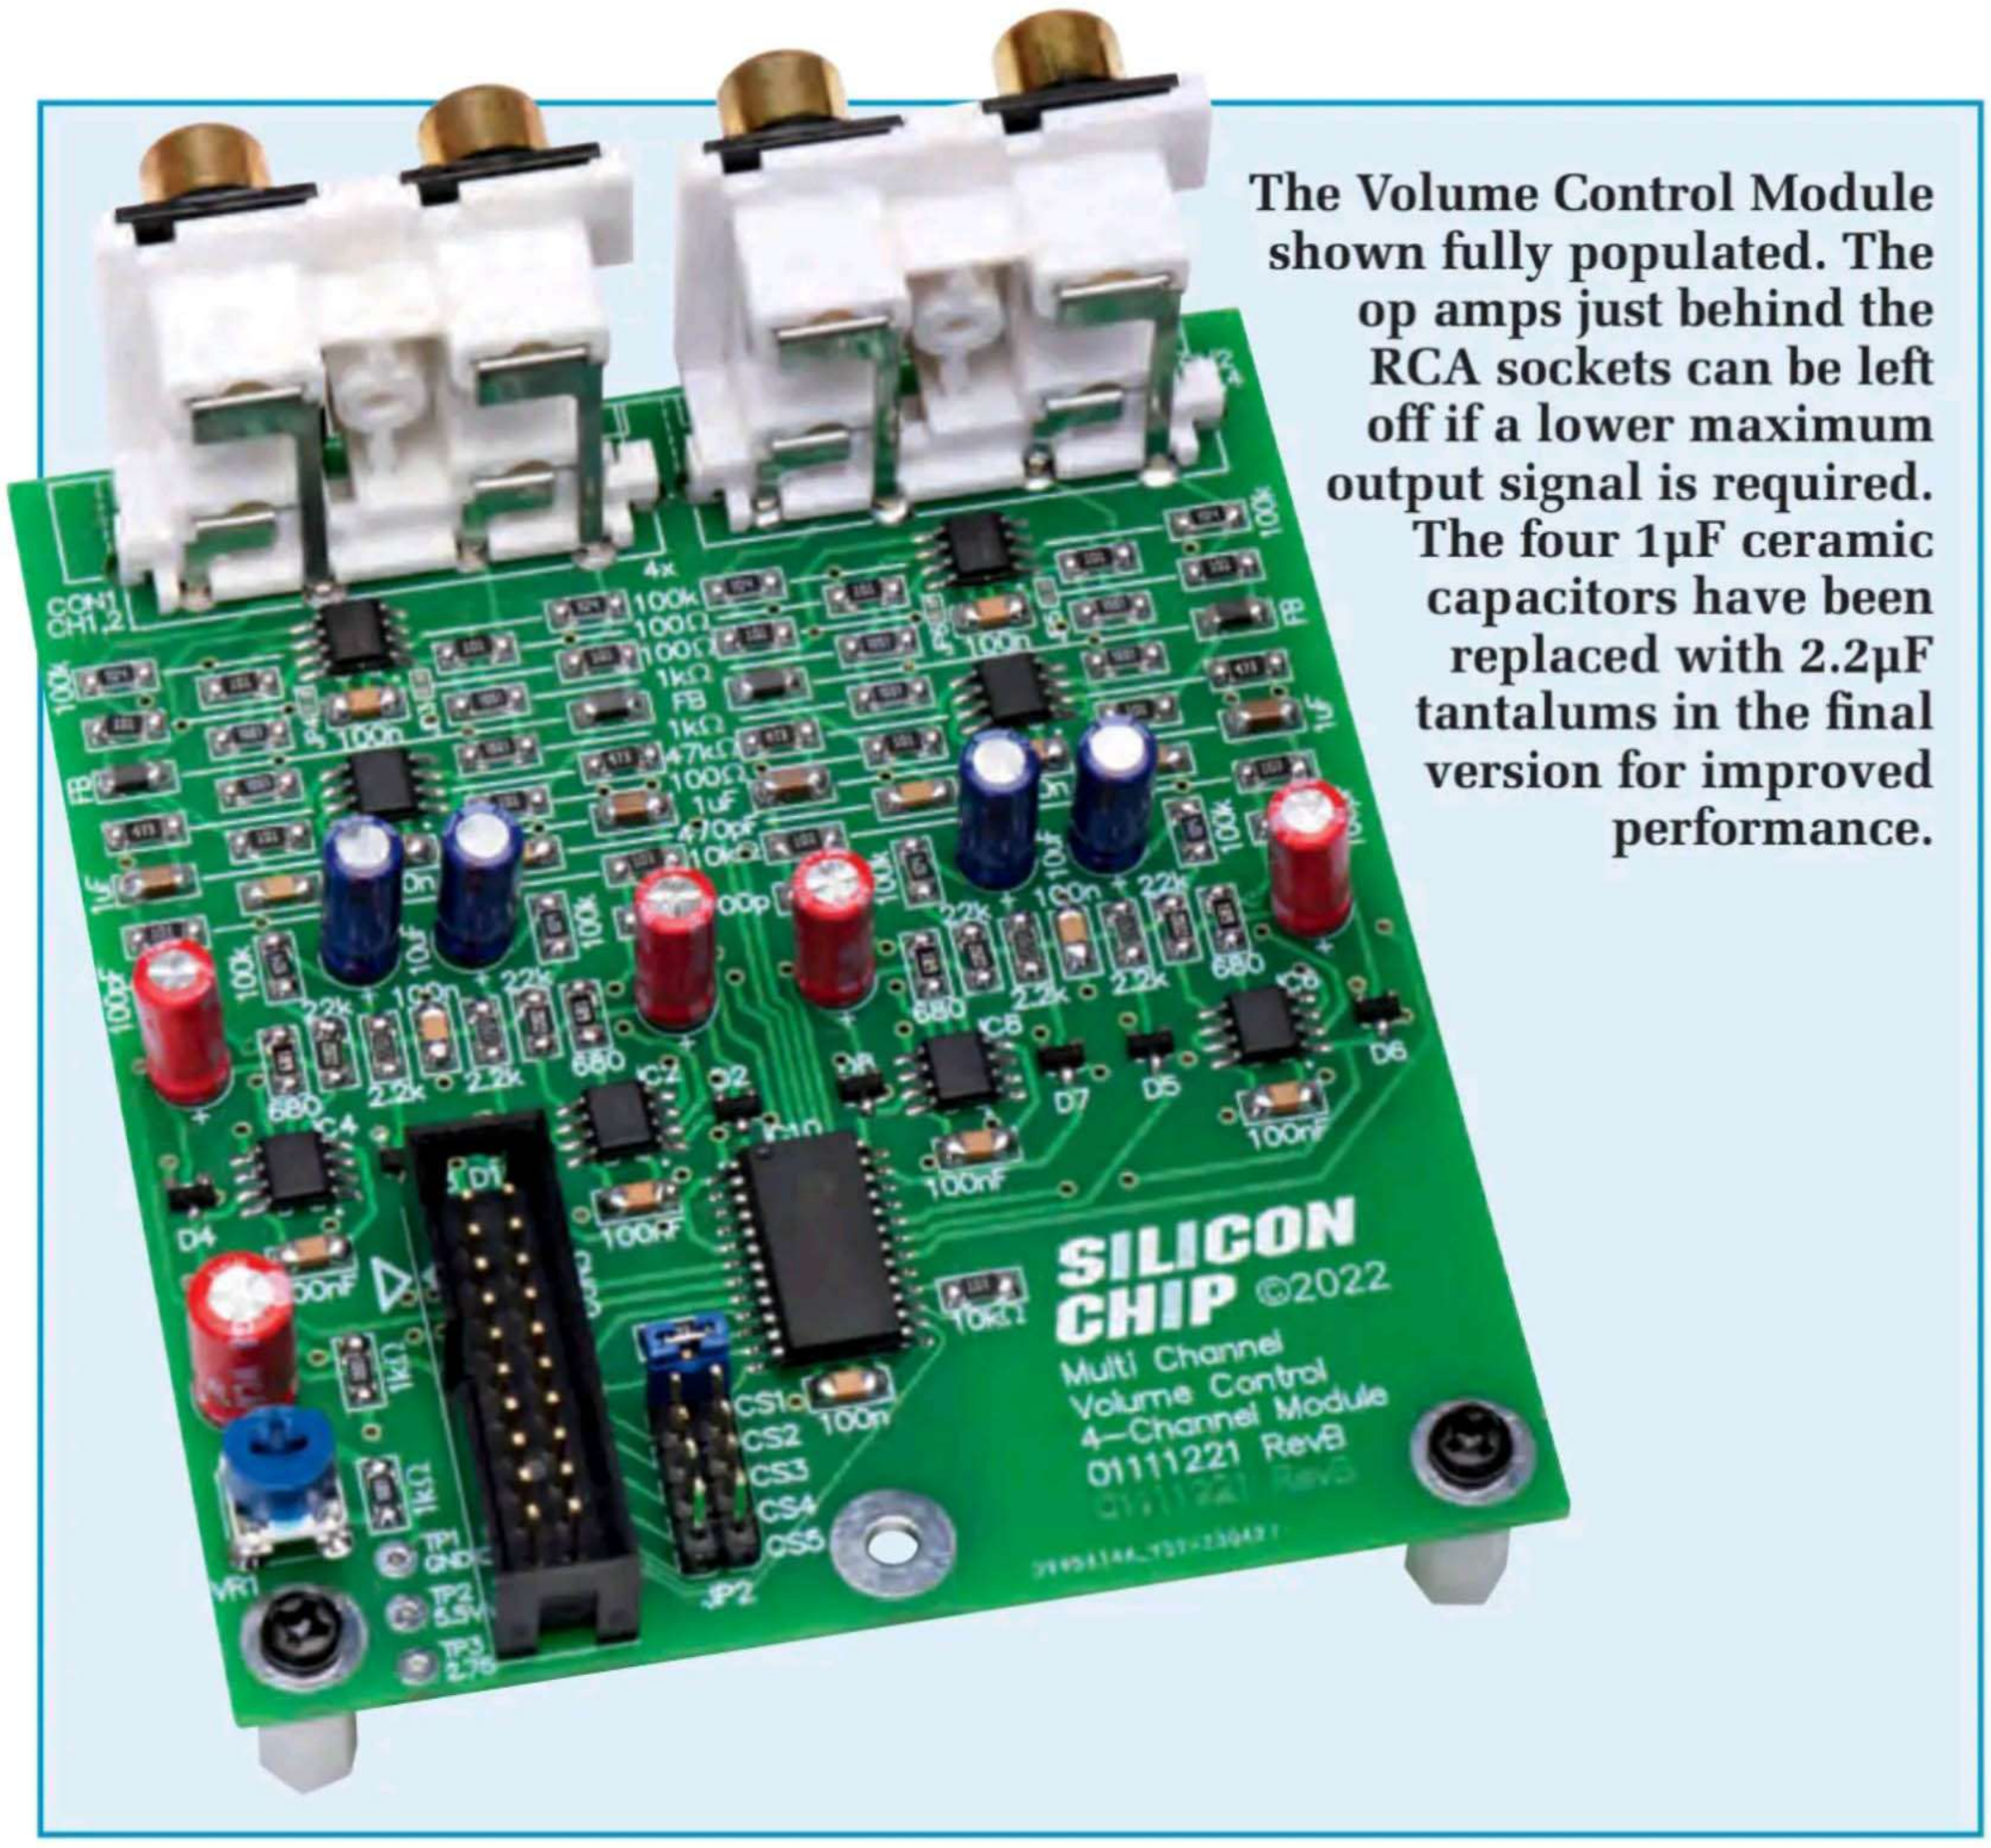

The Volume Control Module shown fully populated. The op amps just behind the RCA sockets can be left off if a lower maximum output signal is required. The four 1µF ceramic capacitors have been replaced with 2.2µF tantalums in the final version for improved performance.

the second with a link across CS2 etc. If you're not sure, install the double-row pin header and place the links as described for testing.

Adjust the 500Ω VR1 trimpot to near its midpoint, then solder it in place (or centre it after soldering).

Next, fit box header CON5. Its key should be to the left, with pin 1 towards the middle of the PCB, as indicated by the arrow on the silkscreen. You could use a double-row pin header at a pinch, although that won't guarantee the correct plug orientation.

Next, mount the nine electrolytic capacitors. Watch out for the polarities (the longer lead is positive, while the stripe indicates negative) and install them as shown.

The last parts to be soldered are the RCA sockets. Their pins and alignment pegs take a bit of wrangling, so ensure their bases are flush against the PCB before soldering them in place.

We also suggest adding a tapped spacer to each of the bottom corners of the PCB. Secure them from above with machine screws. These are used to mount these boards to your choice of enclosure but will also keep the PCBs off your bench during testing.

OLED Module

The optional OLED Module is the smallest of the three. It is little more than a microcontroller, a rotary encoder and an OLED screen. All the components are fitted to one side of the PCB; the other side forms its front panel. You can see this in the Fig.10 overlay diagram and the photos.

Fit the PIC16F15224 microcontroller (IC11) first. Add flux to the PCB pads, rest the micro in place, tack one lead and check its alignment before soldering the remaining leads.

There are four 100nF capacitors and four 10kΩ resistors. None of these are polarised, and can be soldered next. At this stage, you should also add a solder bridge to the CS5 (bottom-most) position of JP7.

Now clean off the excess flux and allow the board to dry. Inspect the solder joints of the smaller components and rectify any concerns. This will be easier before the larger parts are fitted.

If IC11 is not programmed, you will need to fit CON13, the ICSP header. As you can see from the photo, we used a right-angled header fitted as a surface-mounted part.

To program IC11, set your programmer to provide 3.3V, connect it to the ICSP header and upload the 0111122C.HEX file.

Next, solder CON12, the surface-mounting box header. Note the pin 1 marking indicating the orientation. The key for the tab on the cable should face towards the top edge of the PCB.

You could use a standard surface-mounting dual-row pin header if you don't have a box header, but it will lack the keying that ensures the plug is always inserted correctly.

Like any other part, apply some flux, rest the header in place and tack one lead. Adjust the position if necessary, then solder the remaining leads. Since these larger pins are at 0.1in (2.54mm) spacing, you can be pretty generous with the solder. You should be able to look at the gap between the PCB and the box to see that there are no bridges.

Fit the rotary encoder (RE1) next. Mount the encoder using the supplied nut and then add short lengths of component leads to make the connections to the pads below.

When fitting the encoder, ensure that the pins match the PCB silkscreen markings (two pins on one side and three on the other).

Once the encoder is aligned, you can mechanically secure it using the pads on each side of the body. We used lead offcuts around 1cm long, bent about 3mm from one end. Tin the PCB pads and the ends of the leads and then solder the short end of the leads to the PCB.

We used tweezers to hold the other end of the leads while soldering them, then gently bent the other ends of the leads against the pins of the rotary encoder and soldered them together. You can see this in the photo below.

Similarly, the OLED uses short lead offcuts for its four electrical connections. Don't fit the headers to the OLED,

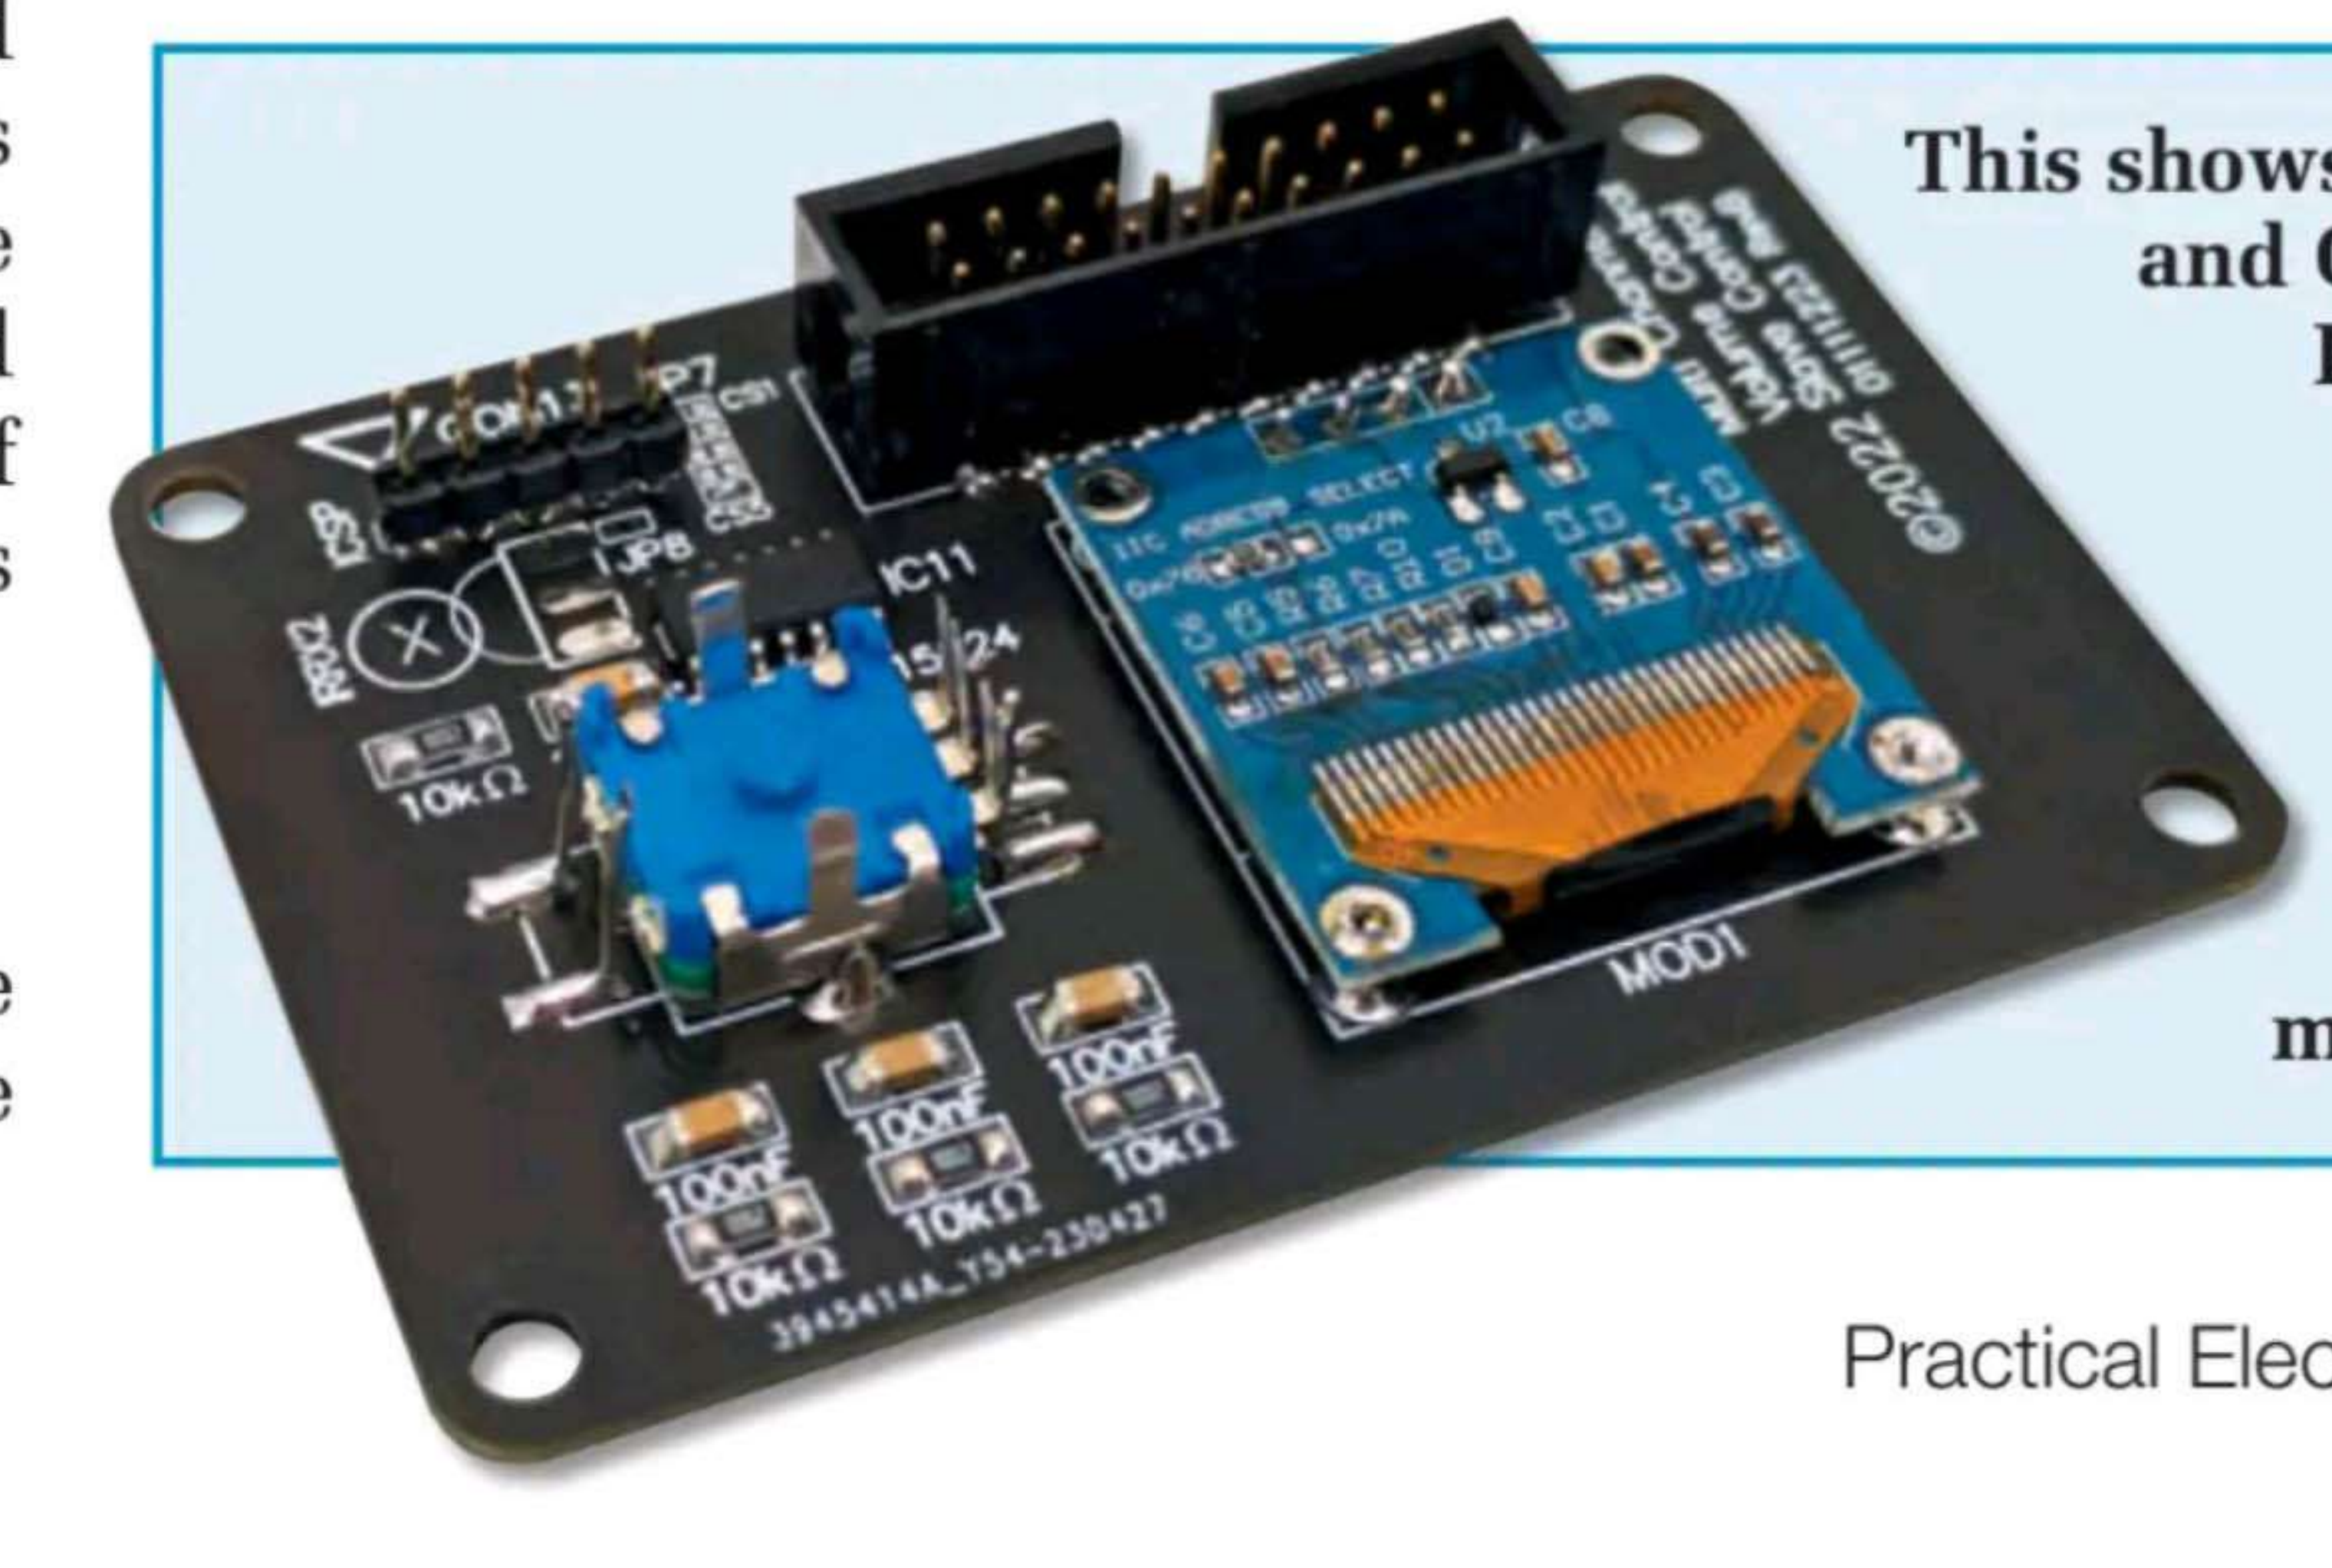

This shows how the rotary encoder and OLED are attached to the PCB. They both use short lengths of wire, such as component lead offcuts, to connect to the PCB. Note how we've soldered a header to CON13 to program the microcontroller in circuit.

Multi-channel Volume Control, part 2

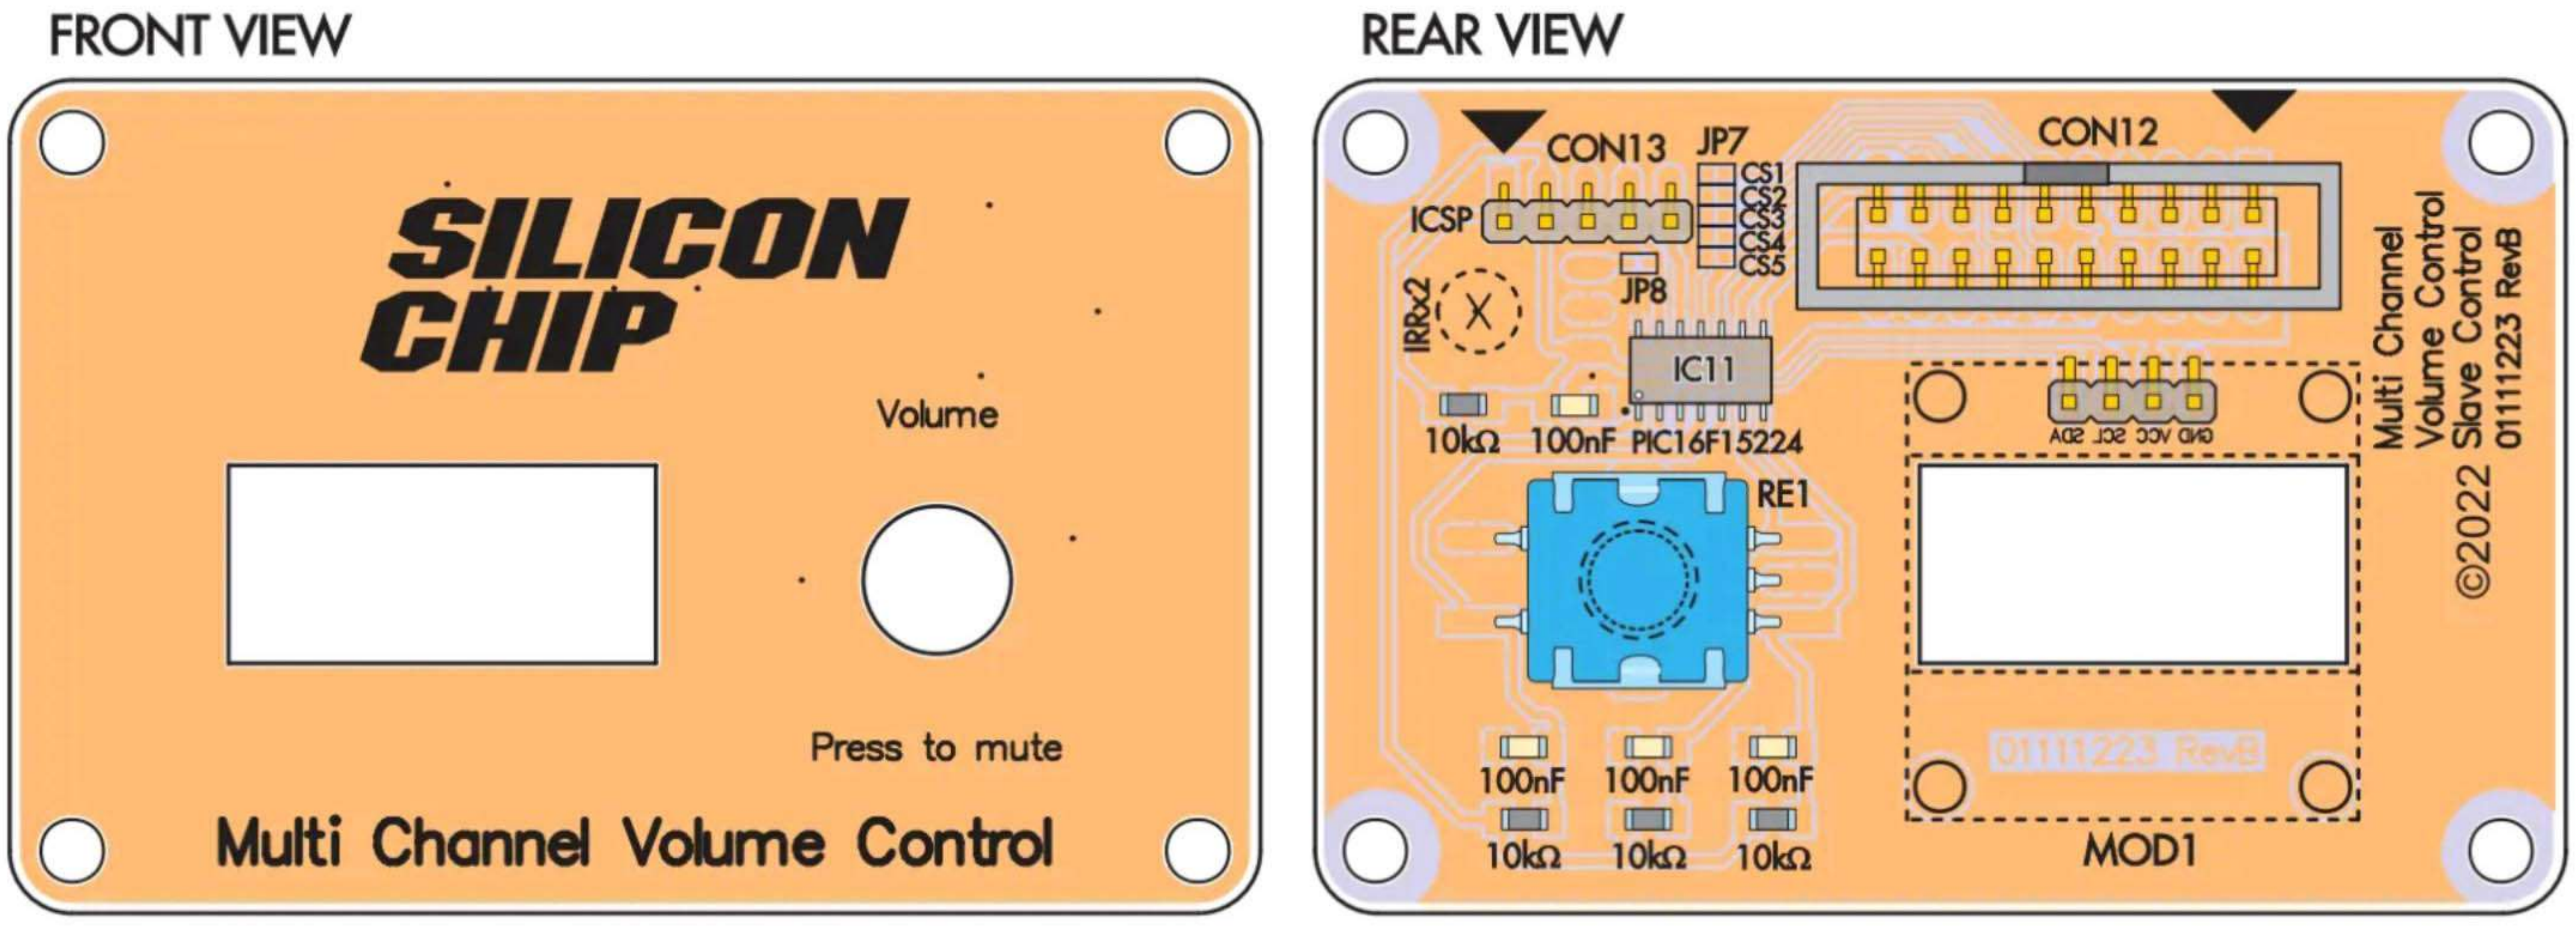

Fig.10: to allow the PCB of the OLED Module to be used as the front panel, all the components are surface-mounted, including the usually through-hole parts. You can also see this in our photos.

as we aren't using them. If one is already fitted, desolder it and clear the pad holes of solder.

Tin each of the four pads on the PCB and then solder a lead offcut vertically. Remove the protective film from the OLED and ease the OLED module down over the leads until it is flush against the PCB.

Gently adjust the position of the OLED so that it is square within the markings on the PCB, then solder each of the four wires to the pads on the Module. Add two more lead offcuts to the two large bottom holes of the OLED and solder them to the PCB pads below.

The OLED should light up if you apply 3.3V and GND (via the ICSP header or pins 9 and 20 of CON12). That's about as much testing as is possible for now.

Ribbon cables

Now we must join all the modules with a custom 20-way ribbon cable. The exact arrangement depends on how you plan to arrange your modules within your enclosure, so we don't have a specific assembly diagram of such a cable.

Fig.11 shows how a typical cable might look. You should use a single length of 20-way ribbon cable and fit it with one 20-way IDC plug along its length for each module you have built. They don't have to be in a specific order, as it is all a single bus.

Pin 1 of each plug must align with the marked pin 1 of the cable (usually red). Otherwise, it doesn't matter too much. The sockets can sit above or below the cable; the endmost sockets should have the cable looped back through their locking tabs to secure them.

It's best to use a designated IDC crimping tool such as Altronics' T1540, but it is possible to use a bench vise with some care. Keep the cable square to the headers and use some pieces of timber on the faces of the vise spread the load. Proceed carefully to avoid cracking the IDC plugs.

Commissioning

Now connect all the modules together with your ribbon cable and wire up your AC supply of choice. A single 12V AC source can connect to CON7 or between the GND and one of the AC phases on CON8. For a 24V AC supply, connect its centre tap to the GND of CON8; the outer 12V taps go to the other terminals of CON8 (it doesn't matter which).

Power on the Multi-Channel Volume Control and verify that the LCD panel shows the expected screen. We still need to perform one last setup step for each of the Volume Modules.

Take a multimeter and confirm that there is 5.5V between TP1 (GND) and TP2 (5.5V) of each Volume Module. If so, adjust trimpot VR1 on each to get 2.75V at TP3. This completes the hardware setup.

We'll now delve into the firmware settings to complete the configuration and then work through the operation of the controls.

OLED Module

If you have an OLED Module fitted, you should be able to operate its controls and see that both displays update together. Screen 2 shows a typical OLED Module display. There is no configuration needed for the OLED Module.

The OLED Module will show three dashes when powered up until it receives data on the ribbon cable. If the dashes persist for more than a few seconds, the OLED module may not be receiving data correctly. In that case, check the ribbon cable and connectors, especially that the IDC plugs are fully clamped around the ribbon cable.

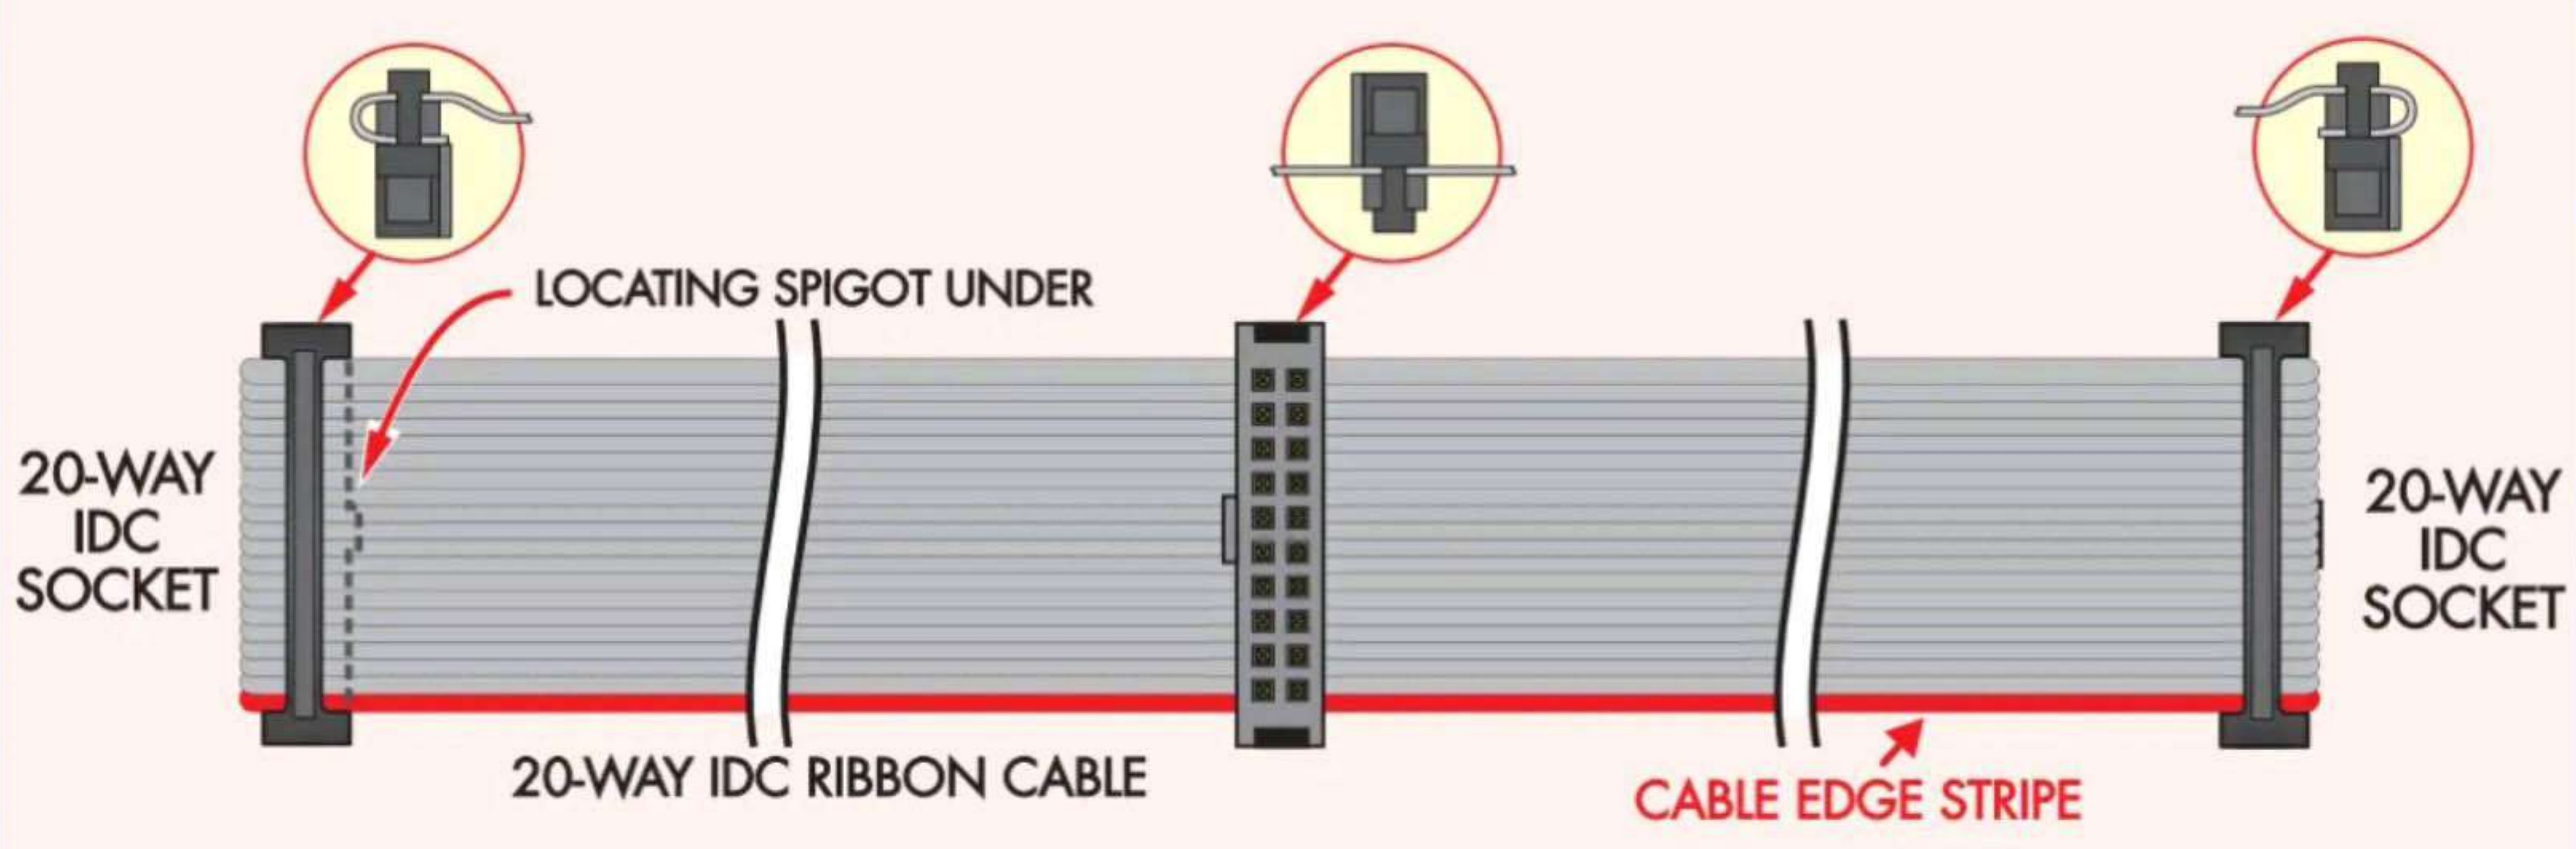

Fig.11: this is only an example of a possible ribbon cable; you might have different requirements depending on your choice of modules. As long as the pin 1 markings align with the same edge of the cable, the cable should work. Note how the keys on the headers on one side of the cable face the same way, opposite to the keys on the other side.

0.0dB
MUTE ON

Screen 2: the OLED Module display will show this on the screen (depending on the MUTE state). If you see three dashes then the OLED Module is not receiving data from the Control Module.

If you find that the operation of the rotary encoder is backward, reverse the connections from the two outer pins to the PCB using short lengths of insulated wire. We haven't seen this happen, but it is an easy fix.

Setup

The default settings for the Multi-Channel Volume Control are to drive 16 channels with an OLED Module connected and the last op amp stages fitted to each Volume Module.

If you have fewer than 16 channels, the 'phantom' channels will not respond, so you won't need to change the settings even if you only have six or eight channels.

The default IR code settings allow the Volume Control to respond to the Jaycar XC3718 IR remote control unit. Use the "−" and "+" buttons to change the volume and the PLAY/PAUSE button to mute and unmute.

The LCD screen should show a red circle when a signal is received. If you don't see a red circle when operating your remote control, its batteries could be flat, or the IR receiver may not be connected correctly. Screen 1 shows the IR telltale.

To enter SETUP on the Control Module, press and hold the SETUP button on the LCD touch panel until the screen changes and you see Screen 3. In general, the ">" button cycles between the different settings, while the "+" and "−" buttons adjust them.

The first four parameters set the IR device code and IR command codes. All the commands must correspond to the same device code. While these can be set manually, the option to 'learn' a code is also available.

Press a button on your transmitter of choice and see that the value in brackets changes; these are the device and command codes the IR receiver detected. You might need to press another button and then your chosen button again to confirm this.

Pressing this area of the screen (around the IR codes) will set the last received device code or command code as the current code. The values are stored in EEPROM and used immediately, so you can easily check that the Volume Control responds to the new IR code as expected.

We have also found a set of codes that can be used with the Altronics A1012A Programmable IR Remote Control. Program the A1012A to use AUX code 0724 (which is for a Yamaha amplifier).

This corresponds to device code 94 and command codes 216 (DOWN), 88 (UP) and 56 (MUTE). You could use the code-learning feature instead of having to enter these manually.

Many other Japanese manufacturers use NEC codes. If the Yamaha code conflicts with existing equipment, a few other codes (from the Altronics A1012A list) that start with 07 also give valid NEC codes that the Volume Control can receive.

The MAX VOLUME setting limits the highest value that the volume can be set to in dB. This can be set as high as 20dB and defaults to 5dB. Disabling the OLED Module is also possible by setting the SLAVE IN USE parameter to 0. If your OLED Module is not responding, check that this is set to 1.

The LEVEL OFFSET parameter provides an adjustment to the overall gain. If you have omitted the last op amp stage on the Volume Modules, set this to -6 to account for the loss of the last ×2 gain stages.

The next parameter changes the number of channels in use; this is the number of channels driven by the Volume Control. This should be a multiple of four and match the number of Volume Modules you have installed. If in doubt, set it to the maximum possible.

Say you have two Volume Modules and are using six channels; in that case, set it to eight to ensure the two spare channels are set to safe levels. It can't be set higher than 16 if the OLED Module is enabled.

For these settings (apart from the IR codes), the values in brackets show the lower and upper limits of what these parameters can be set to.

The remaining settings are offsets (in digital potentiometer steps) that can be applied to each channel. This can be used to adjust the balance between different speakers.

A short press on the SETUP button returns to the normal display. Screen 4 shows what the display looks like when MUTE is active. The EEPROM text is also yellow, indicating that the current state has not been saved to

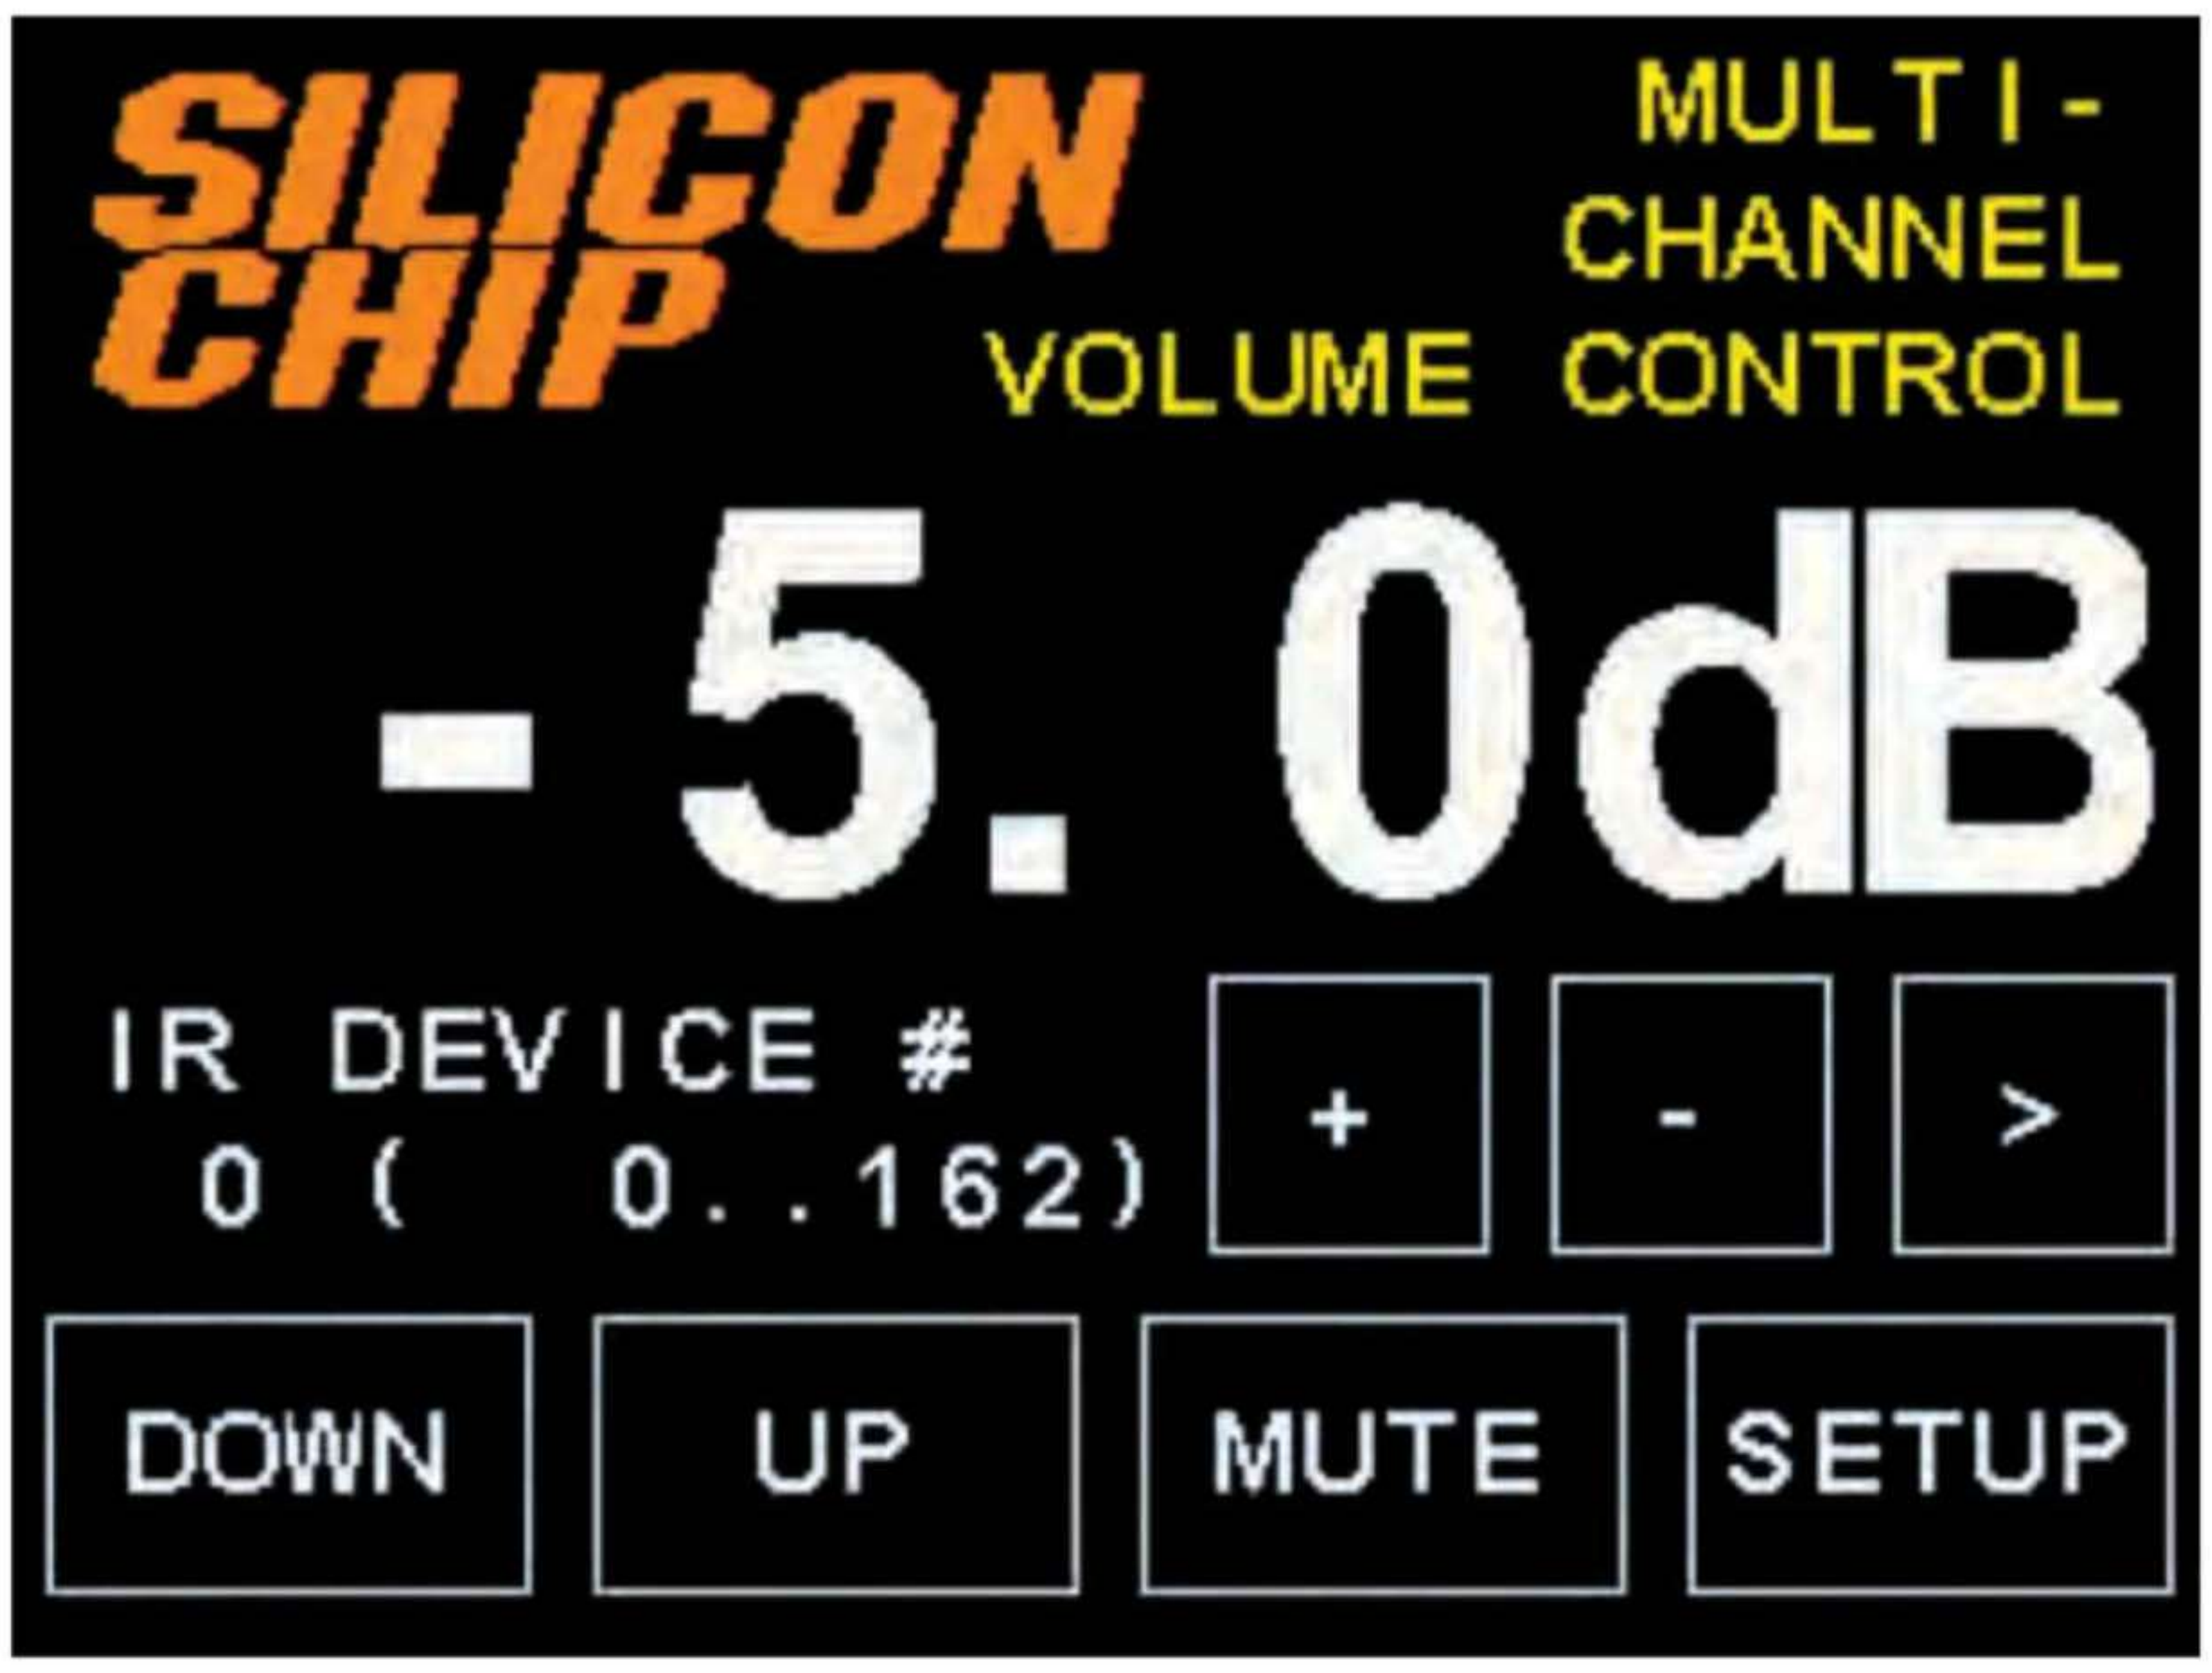

IR DEVICE #
0 (0..162)

Screen 3: during setup, part of the screen is turned over to the setup parameters and buttons. Press and hold the SETUP button for five seconds to get to this screen and start the setup process.

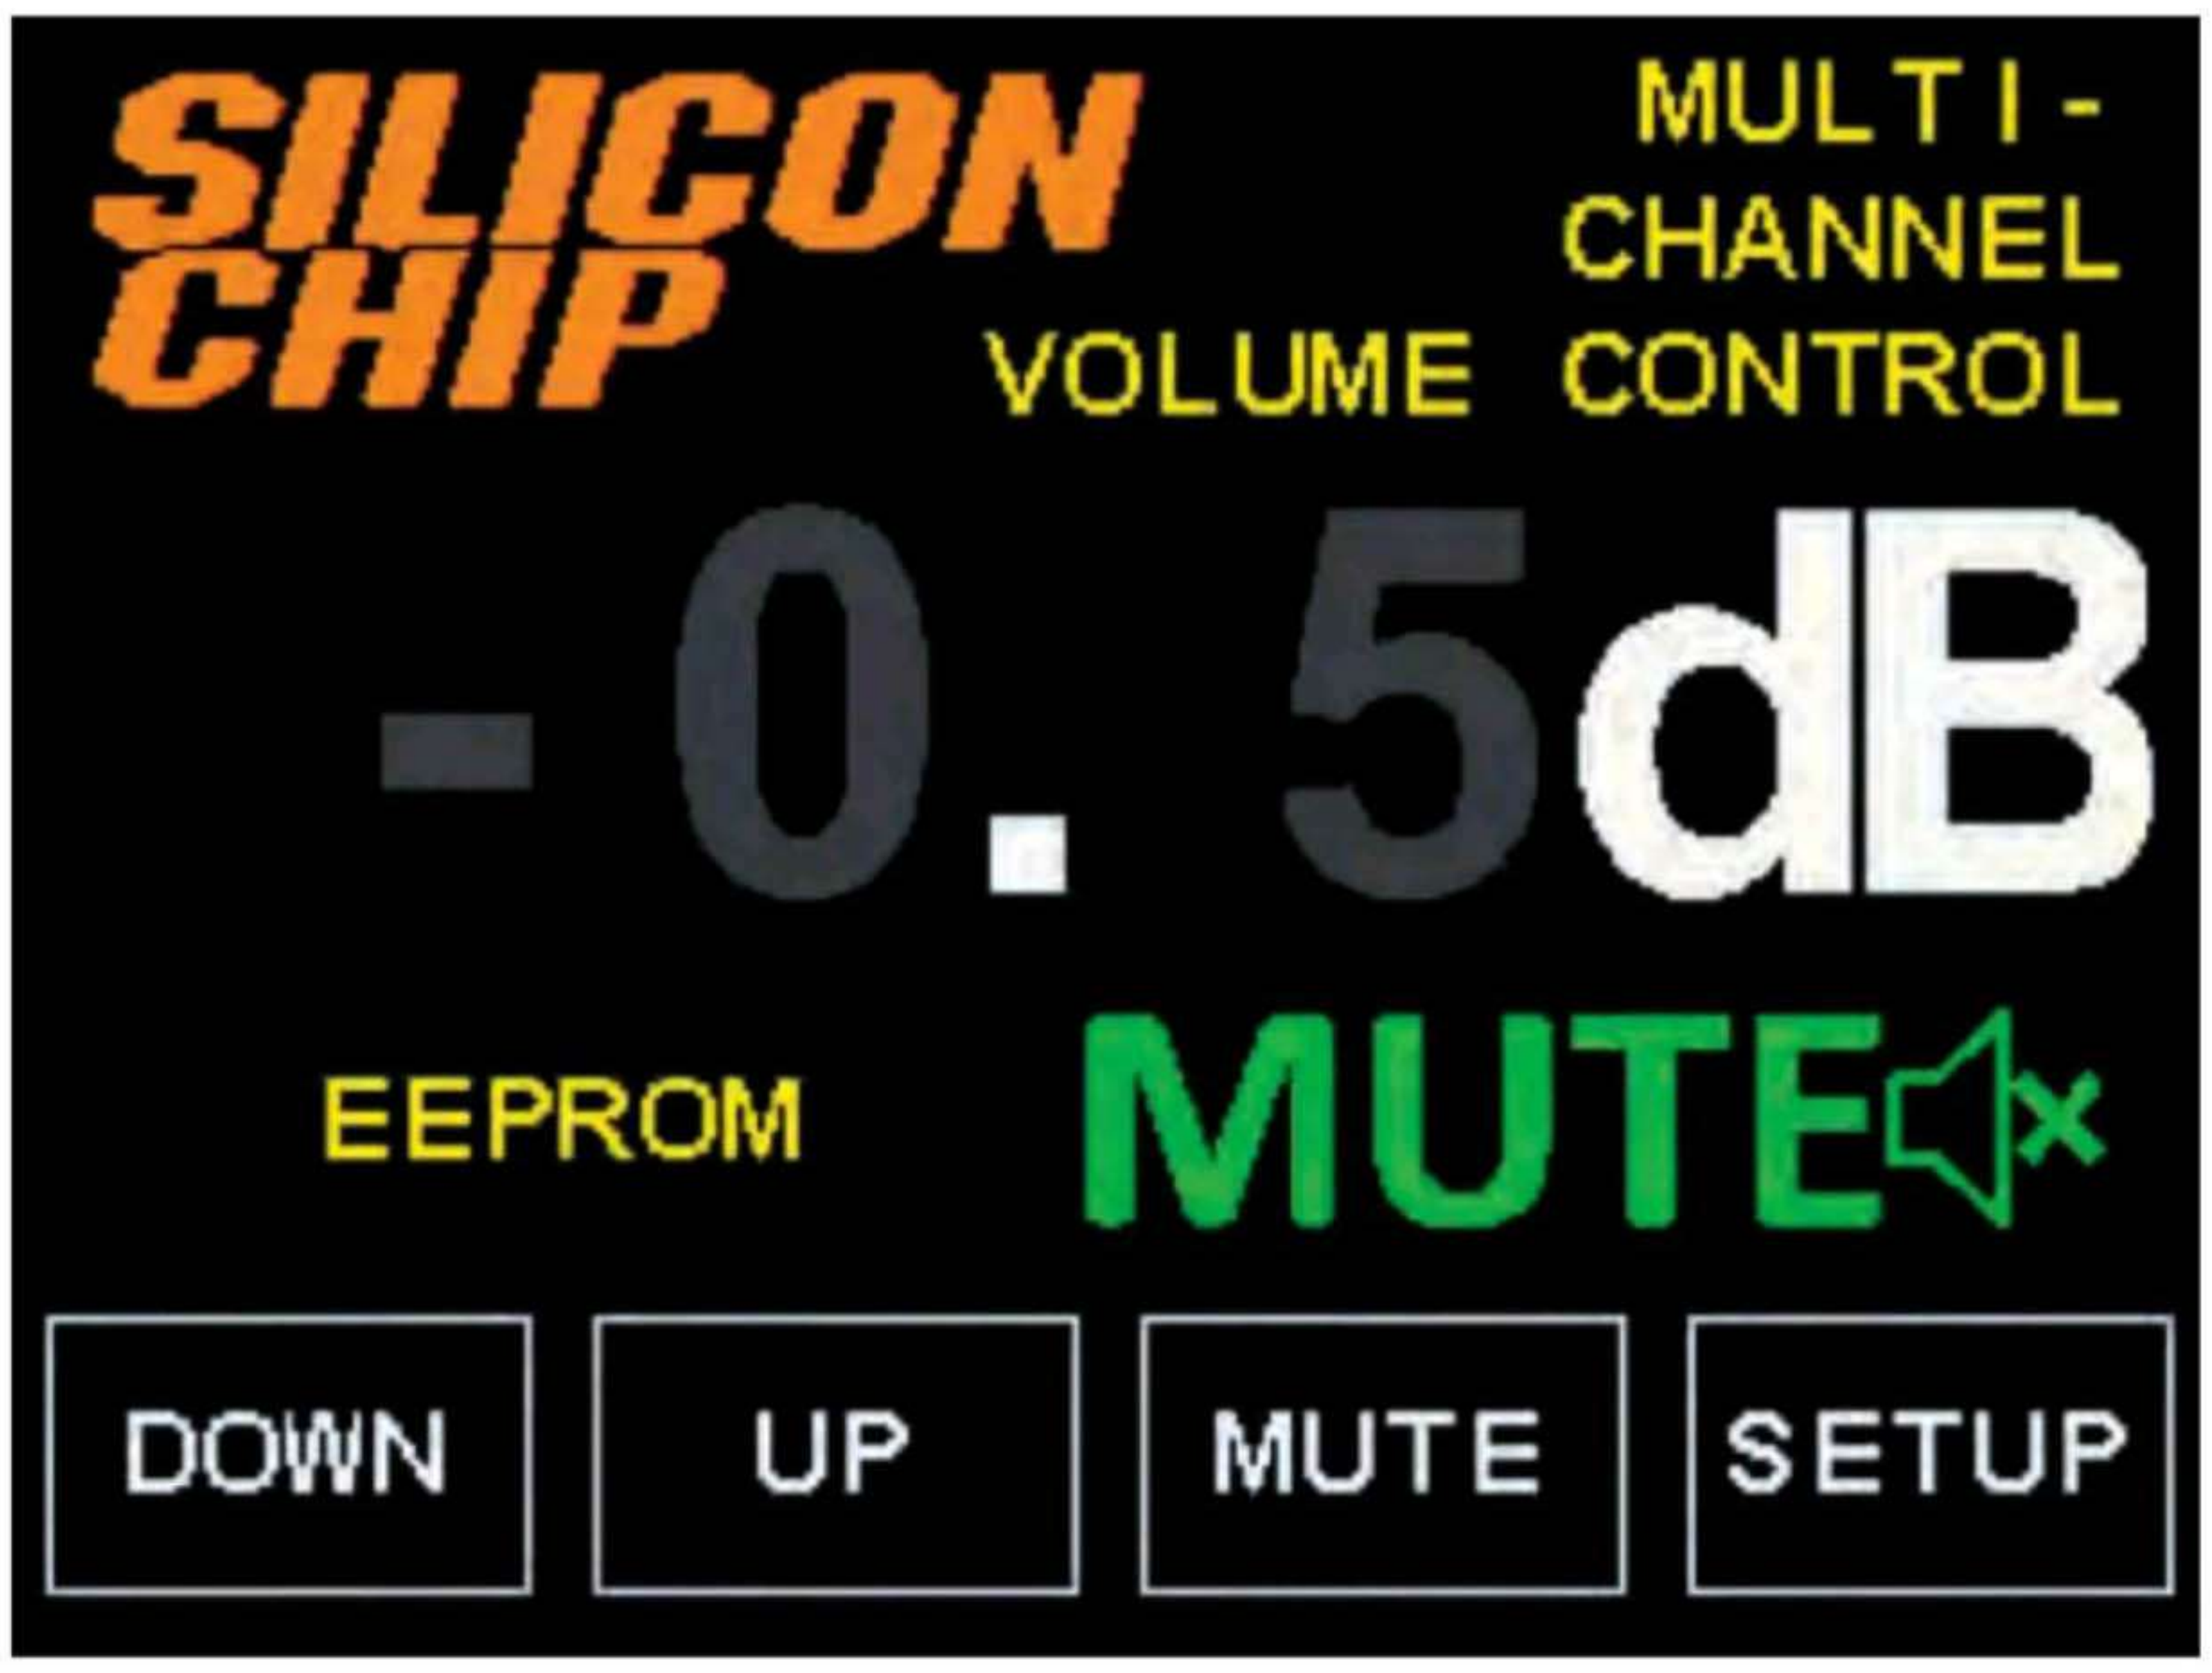

Screen 4: when MUTE is active, the screen changes to look like this. The yellow EEPROM text means that there are unsaved changes. After 10 seconds of no activity, the state (volume and mute) is saved and will be reloaded if the Volume Control is switched off and then on again.

LCD screen backlight modification

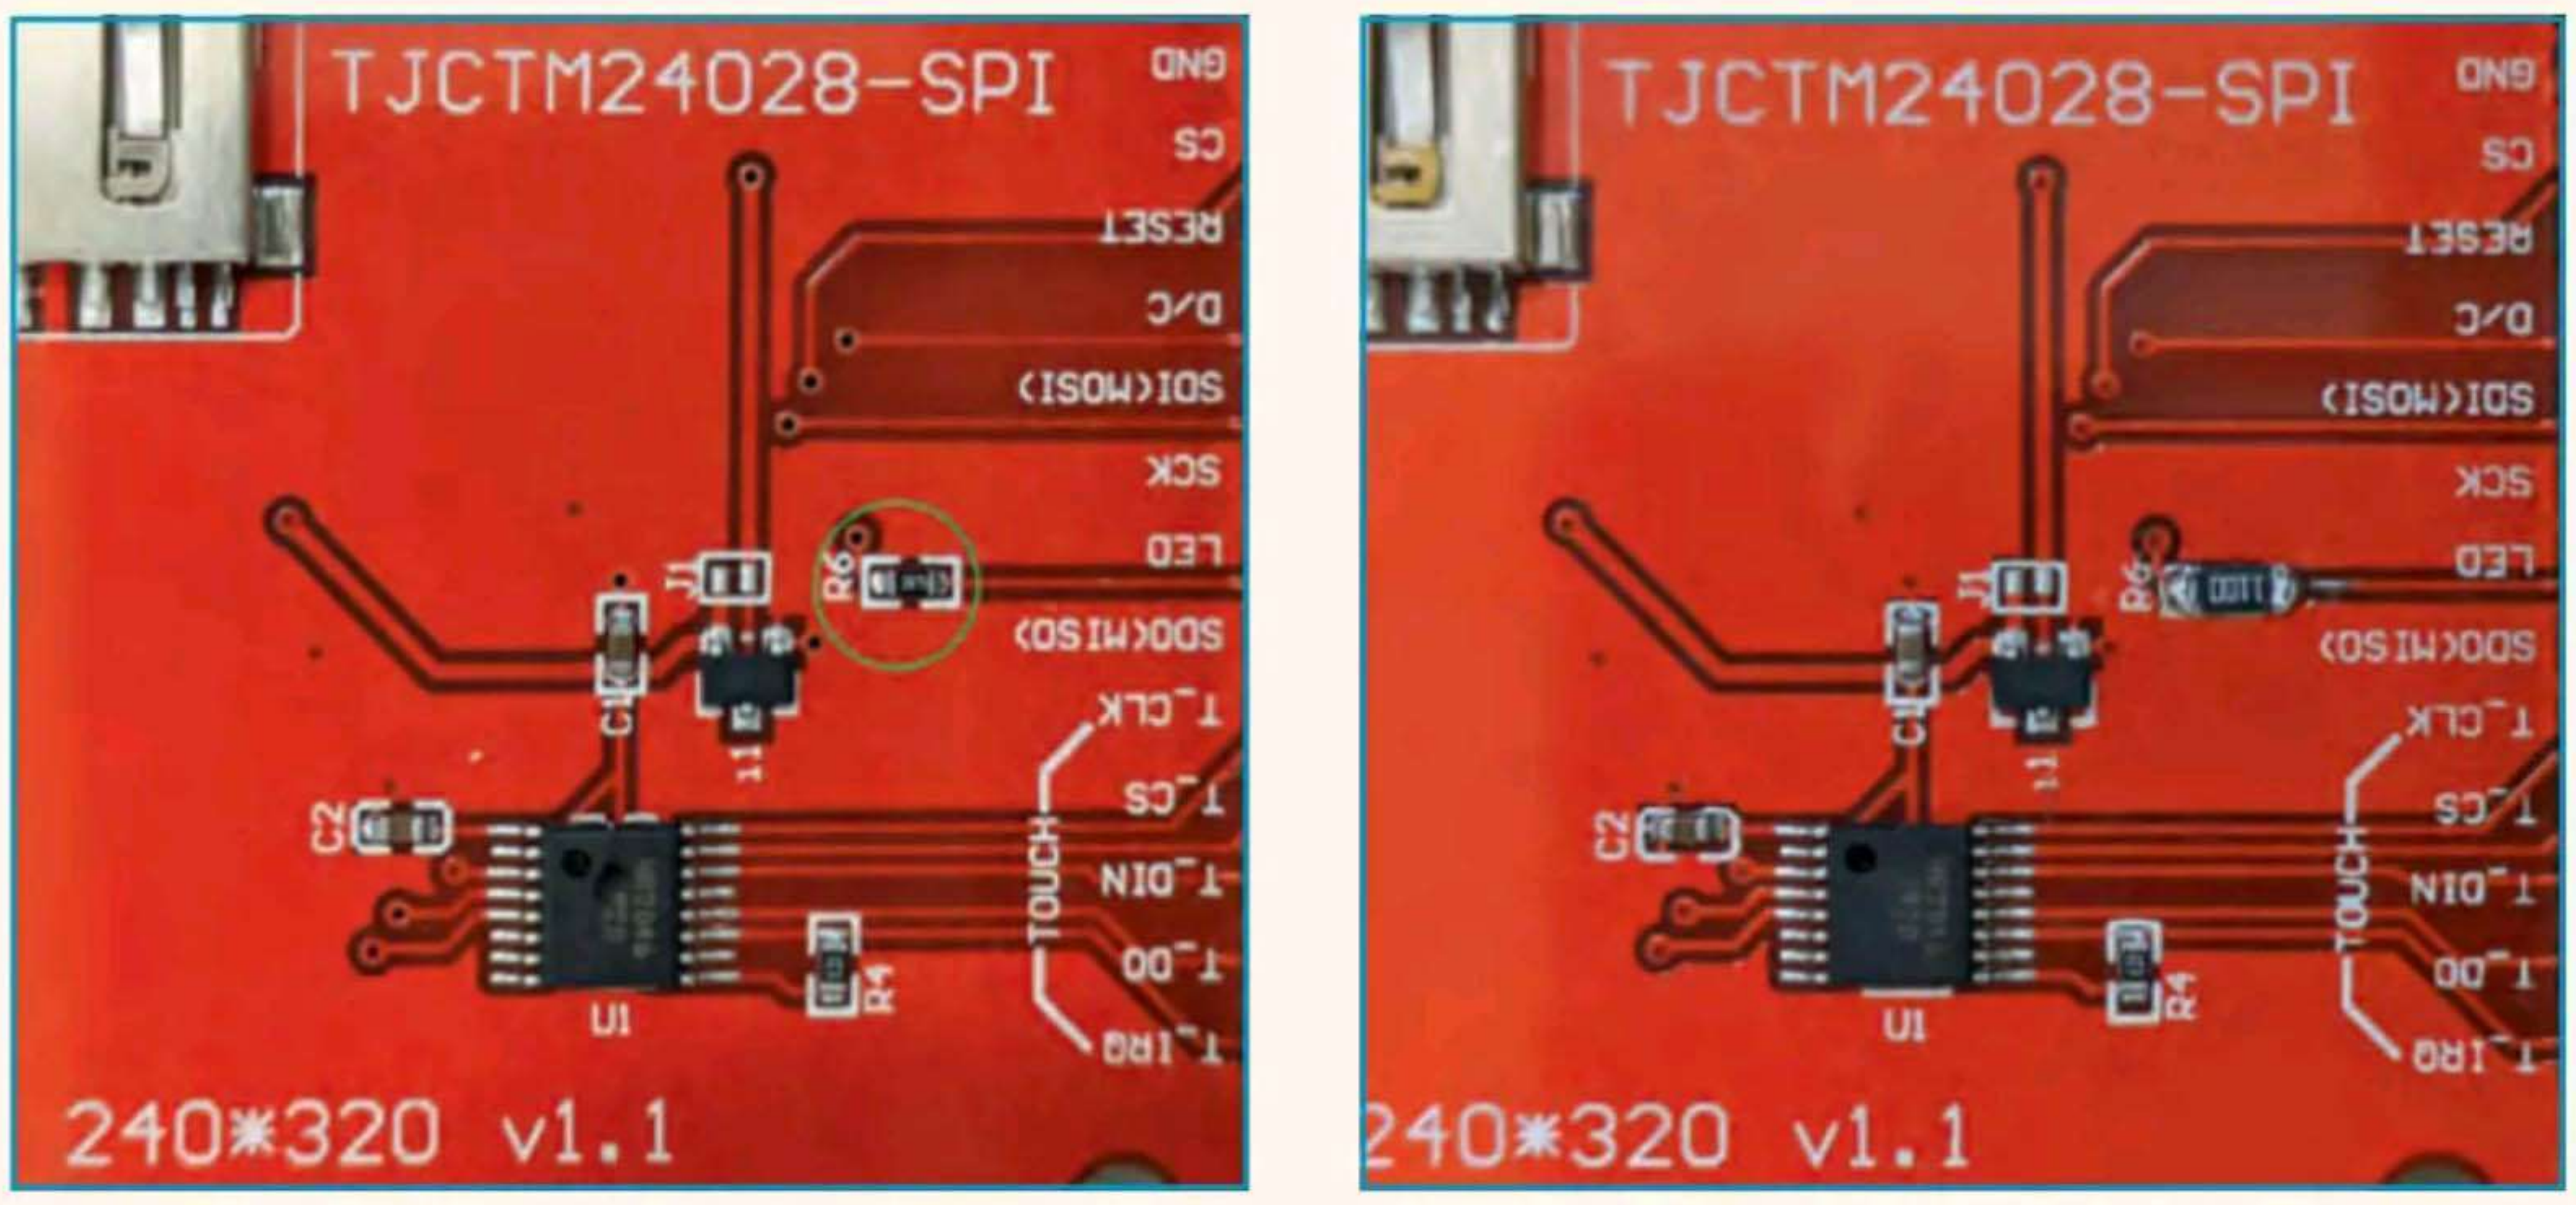

Figs.b & c: a V1.1 touchscreen LCD module before (left) and after (right) replacing R6 with a 110Ω resistor to reduce the backlight current.

There are two common variants of the 2.8in LCD touchscreen panels. The main difference we noted is that the touch panels register differently, requiring different calibrations.

As we mentioned in last month's installment, the Multi-Channel Volume Control is programmed to handle these variations.

Another difference is in the circuitry of the LED driver for the LCD panel backlight. The two variants we have seen are marked v1.1 and v1.2, as shown in our photos (adjacent and below).

Both versions have an XC6206 3.3V regulator to power the LCD controller from the panel's V_{CC} pin and an XPT2046 touch controller IC to provide an interface to the resistive touch panel.

Fig.a highlights how they differ in their connections to the LED control line (one of the pins on the 14-way header). For the v1.1 boards, this line connects directly to the LEDs and then ground via a series ballast resistor.

The later v1.2 boards use the LED control line to drive a low-side NPN S8050 transistor (Q1). The LEDs are wired to the V_{CC} line, so when the transistor switches on, current flows via a ballast resistor and the transistor to ground.

The v1.1 board design lends itself to dimming by an extra series resistor in the LED line. For example, the original Micromite LCD BackPack (described in the May 2017 issue) used a trimpot for manual backlight adjustment. The v1.2 boards do not allow that, so we have tended to use it less and less over time.

While both arrangements can be driven by a high-current PWM signal (which could be provided by Q1 and Q2 of the Control and Power Supply Module), we have avoided using PWM in this project due to the resulting digital noise.

So we looked into how to modify the LCD panel to adjust the backlight current linearly. Fortunately, changing the LED ballast resistor works well enough, which is what we did.

Figs.b & c show a v1.1 board before and after modification. The green circle in Fig.b shows the resistor in question, originally 3.9Ω and designated R6.

The original resistor was an M1608/0603 (1.6 × 0.8mm) part, but we replaced it with a larger M3216/1206 (3.2 × 1.6mm) part by scratching back some of the nearby solder mask to allow the larger part to be soldered.

This must be done carefully as the surrounding copper area is connected to ground, and a bridge here will short the incoming LED signal to ground. We used 110Ω resistors for our tests because we had a few left over from building our prototypes.

Figs.d & e show the v1.2 LCD panels before and after the changes. Here, the resistor is marked R5 and is 8.2Ω. We did the same thing, scraping some of the solder mask back to bare copper before soldering in the replacement resistor.

You might even be able to solder in an axial leaded resistor by bending its leads back until they are nearly touching.

Resistor value

The 110Ω resistors were great at keeping the noise and heat down but the resulting backlight brightness is too dim for a well-lit room. We suggest 22Ω as a good compromise. A 100Ω trimpot in series with a 10Ω resistor would be a good choice if you want to tweak the brightness to suit your specific conditions.

Our Control Module kits will include a 22Ω M1608/0603 SMD resistor so you can make this modification with a direct resistor swap.

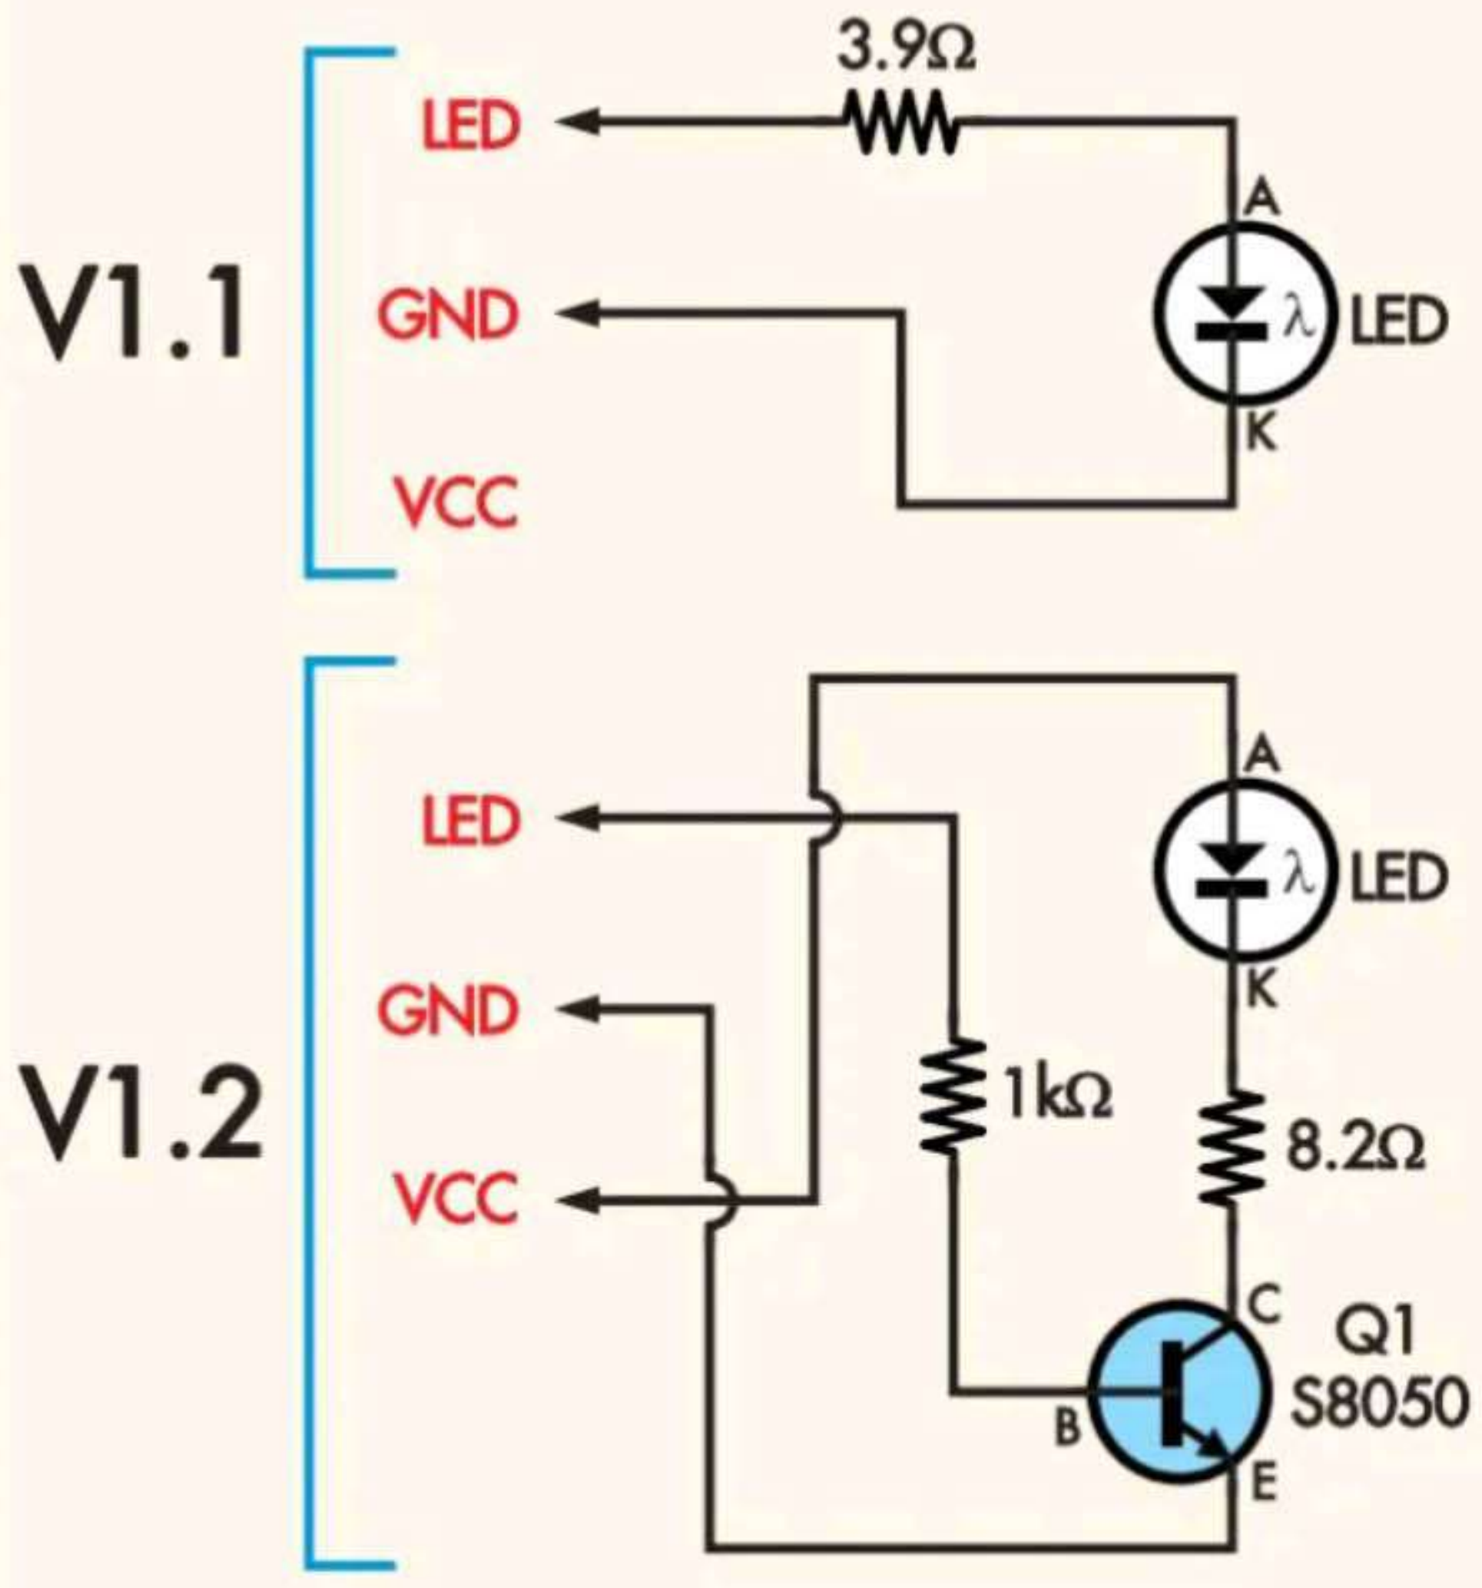

Fig.a: the LED control lines for the V1.1 and V1.2 LCD modules.

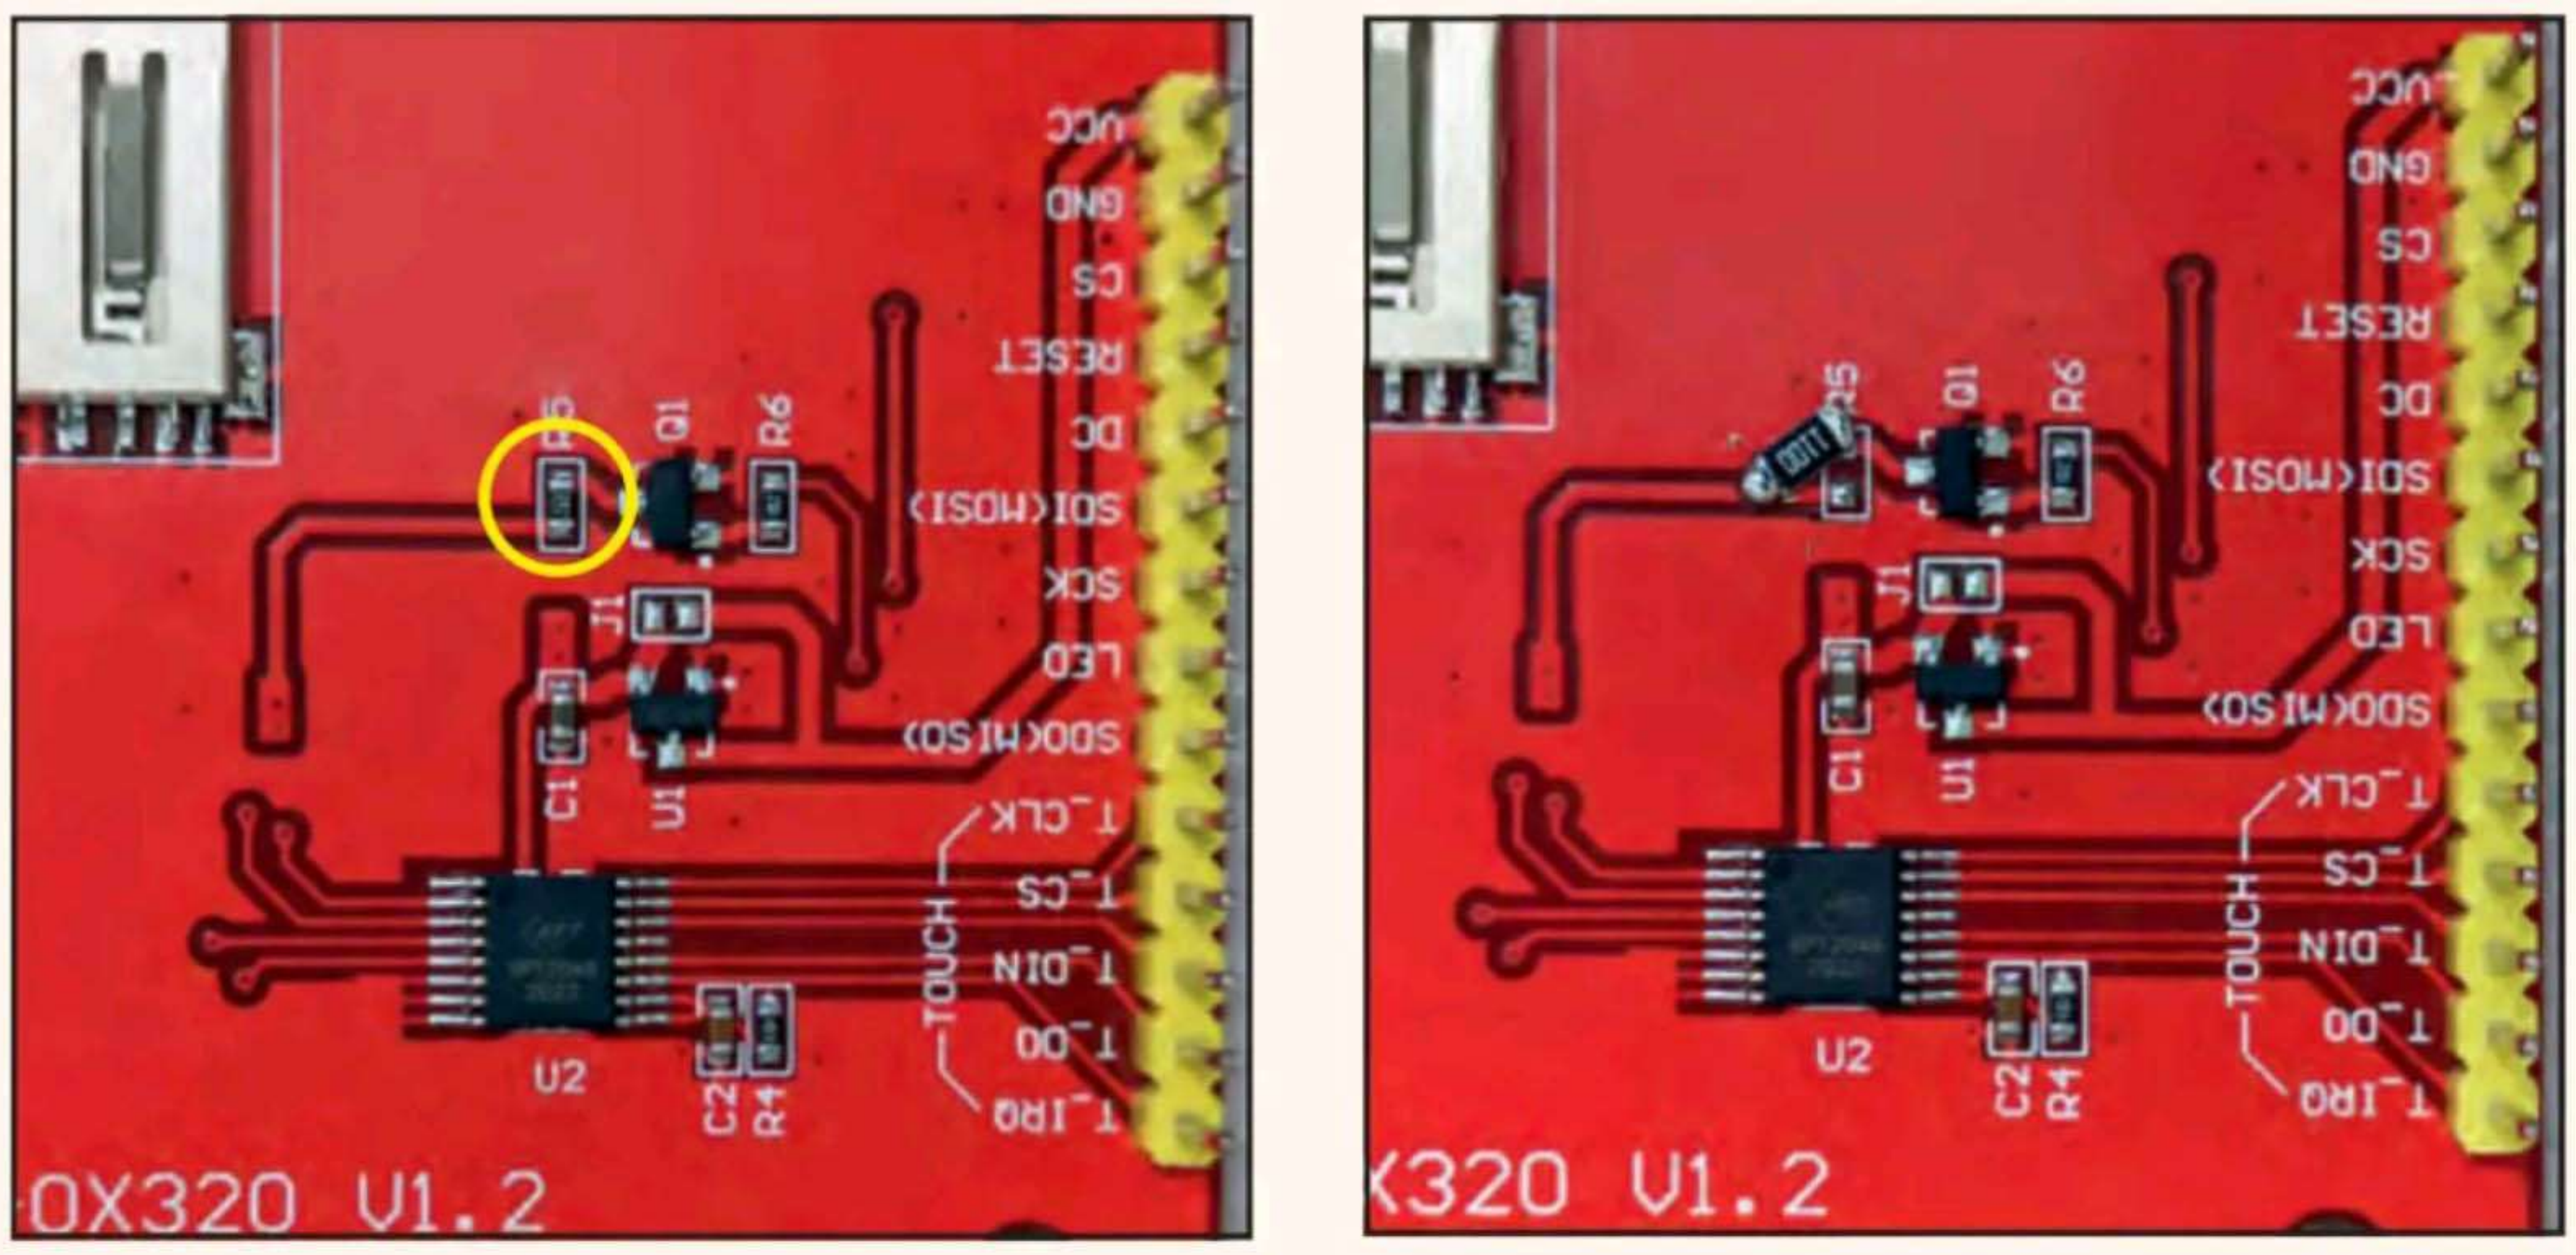

Figs.d & e: for the V1.2 LCD modules, we needed to scrape some of the solder mask so we could fit a larger resistor for R5.

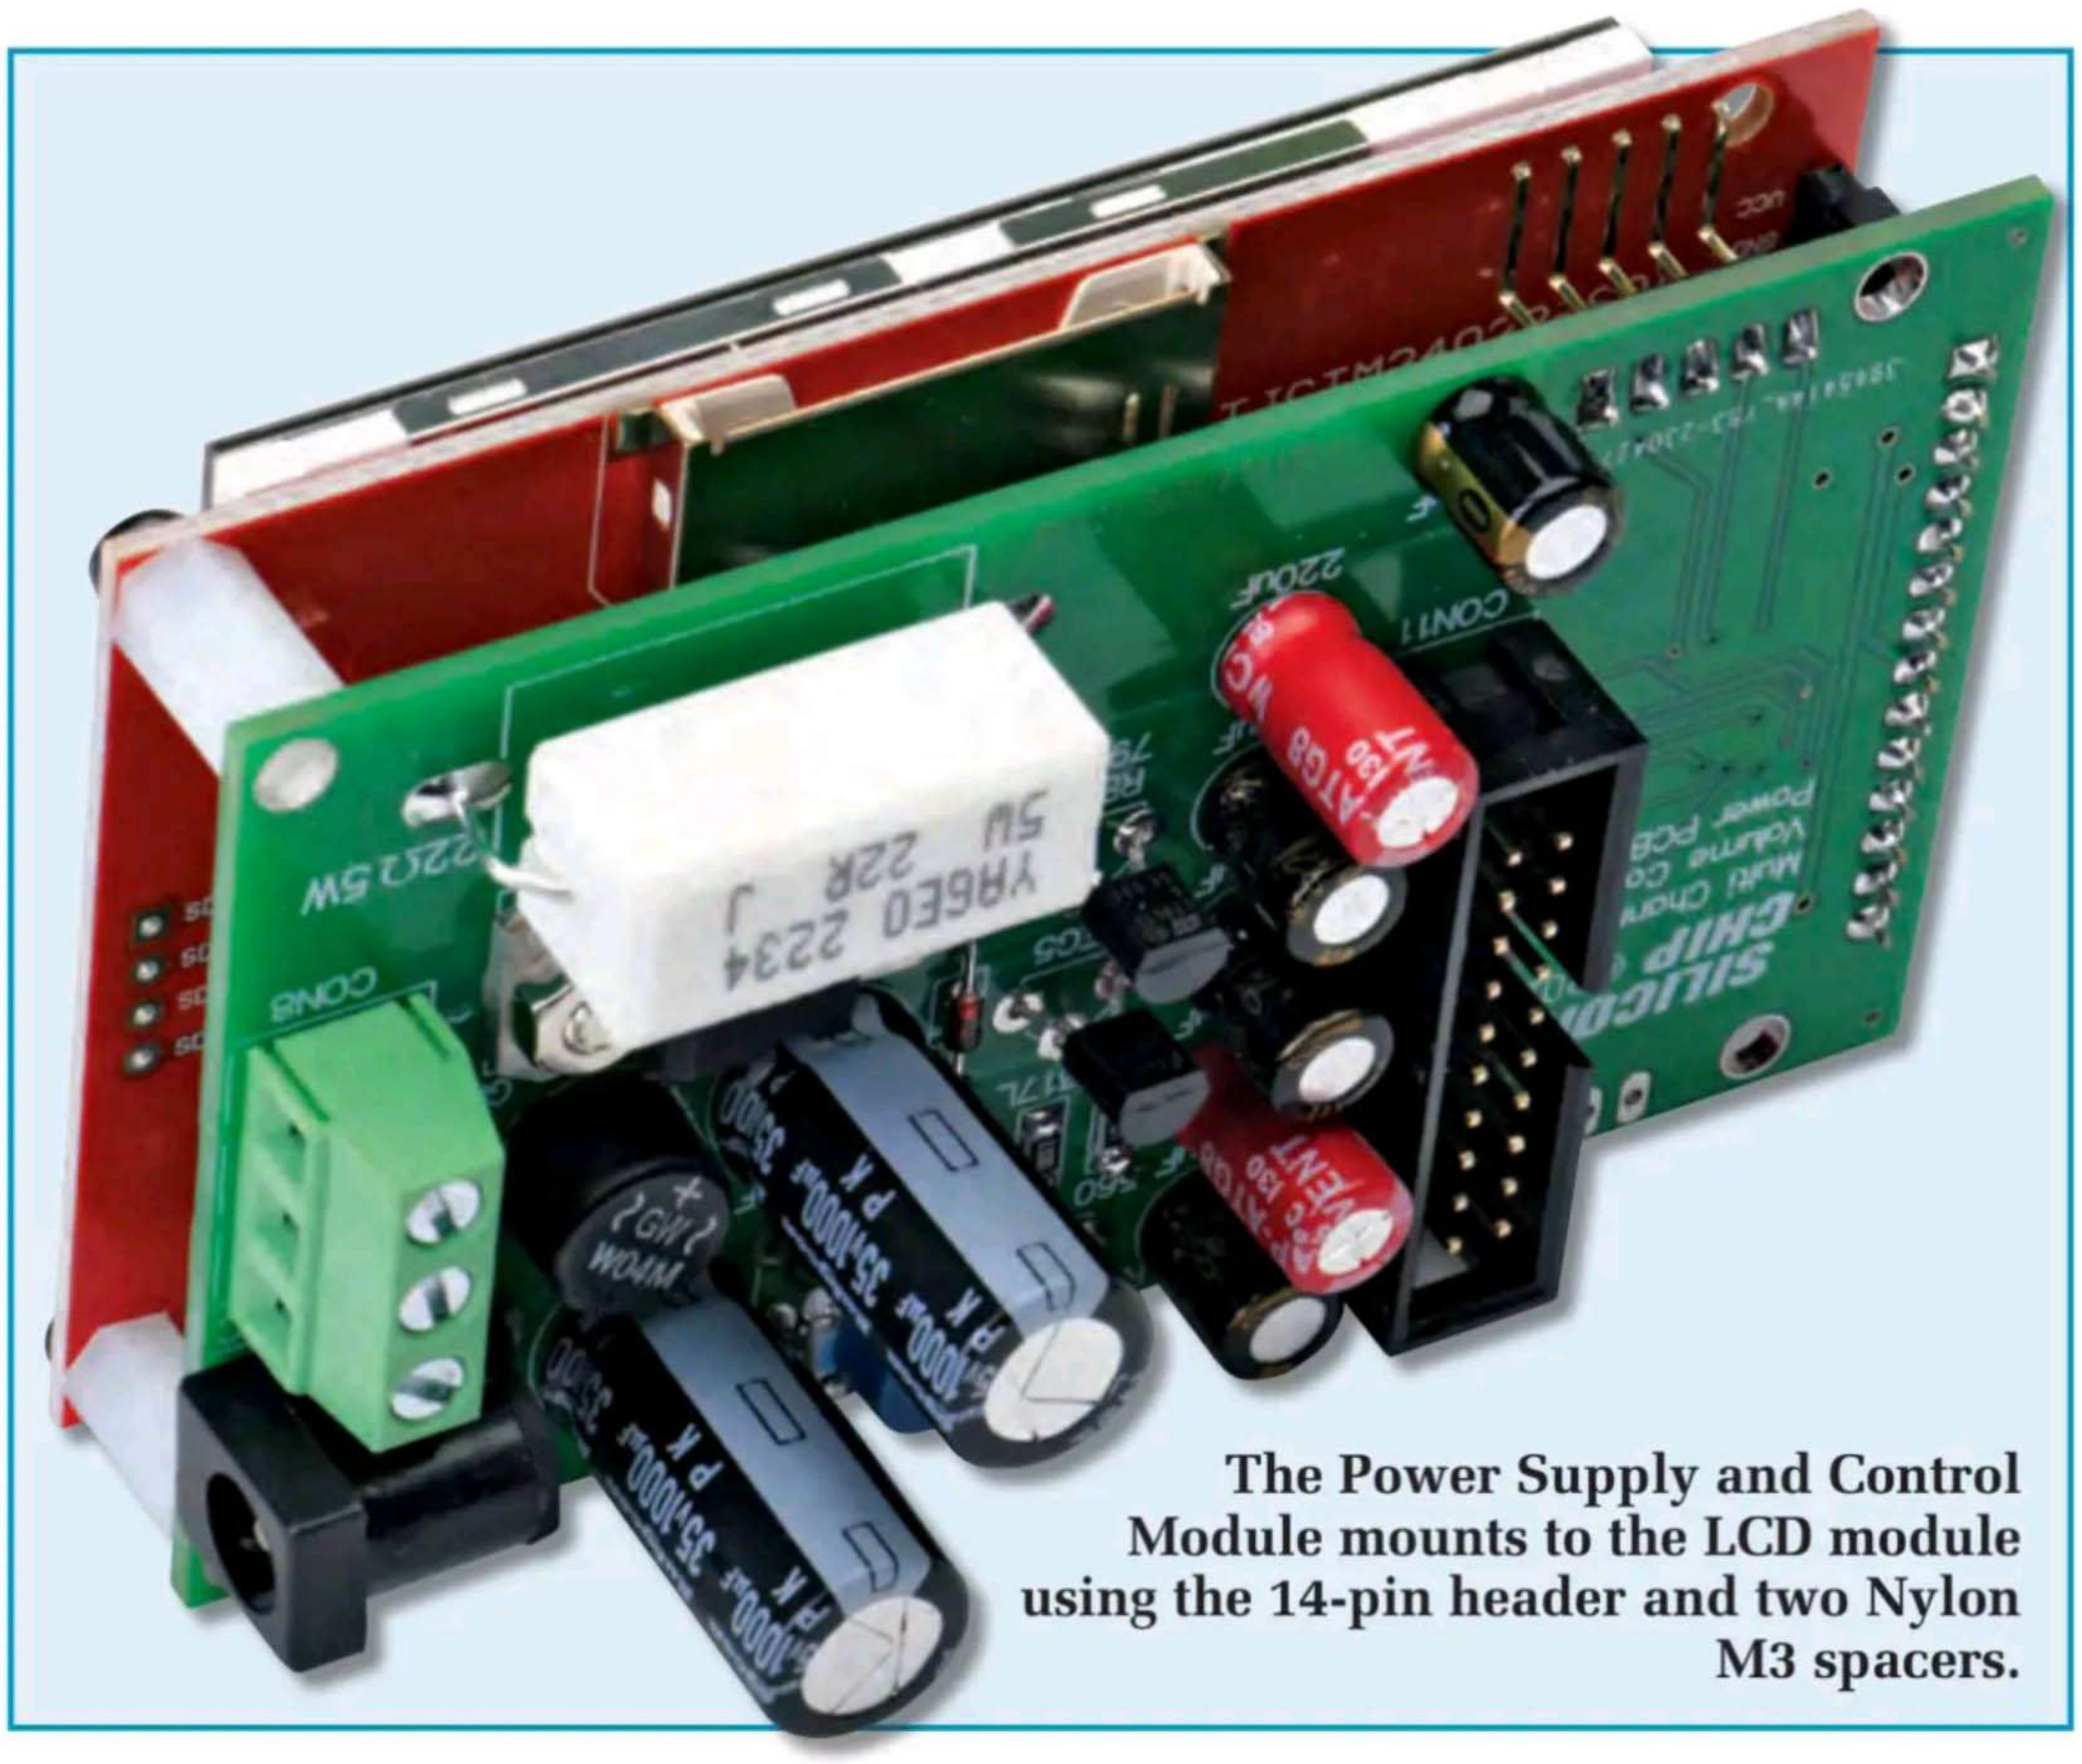

The Power Supply and Control Module mounts to the LCD module using the 14-pin header and two Nylon M3 spacers.

EEPROM; that happens automatically after 10 seconds of no further activity.

Installing the modules

To help you fit the modules into your desired enclosures, Figs.12 & 13 are cutting diagrams of the display cutouts for the Control & Power Supply Module and the OLED Module.

The cutout for the Control & Power Supply Module is essentially the same as for the 2.8in LCD module. You could even consider using a laser-cut acrylic lid from the Silicon Chip online shop, such as SC3456 (siliconchip.au/Shop/19/3456) as a bezel to neatly mount the LCD panel.

This acrylic piece is intended to fit onto a UB3 Jiffy box and is 68mm tall, so it will be too tall for a 3U rack unit. Otherwise, refer to Fig.12 for the dimensions of the square cutout and screw holes to suit the 2.8in LCD.

You will also need to create a hole for the IR receiver if you are using it. Its exact position depends on how you have fitted it.

If mounting the LCD inside a metal enclosure, we recommend using a plastic bezel or foam tape to prevent the LCD pins from shorting against anything.

Fig.13 shows the outline for the OLED Module. The outermost dimensions (76.5 × 51mm) are the outline of the Module, so you can start by marking these onto your enclosure. Use something erasable or work inside the enclosure, as these will be visible once the Module is fitted.

Now add another set of lines 4mm inside these and yet another set of lines 4mm inside these; thus, the second set of lines is 8mm inside the Module's border. These twelve lines will allow you to drill four holes and cut out the panel, as shown in Fig.13. Note that the inner cutout area does not need to be precise. You should leave enough material for the screws to hold.

If the panel is metal, it is worth attaching some foam tape around the perimeter at the back, where the OLED Module attaches. This will prevent the case from scraping the solder mask and possibly shorting against the PCB traces.

Completion

The Multi-Channel Volume Control is intended to be a 'subsystem' within a system such as a multi-channel amplifier, so it is up to you how you connect it to your equipment of choice.

As for the RCA sockets, the white upper connections are the inputs, and the red lower connections are the outputs of each Volume Module. **PE**

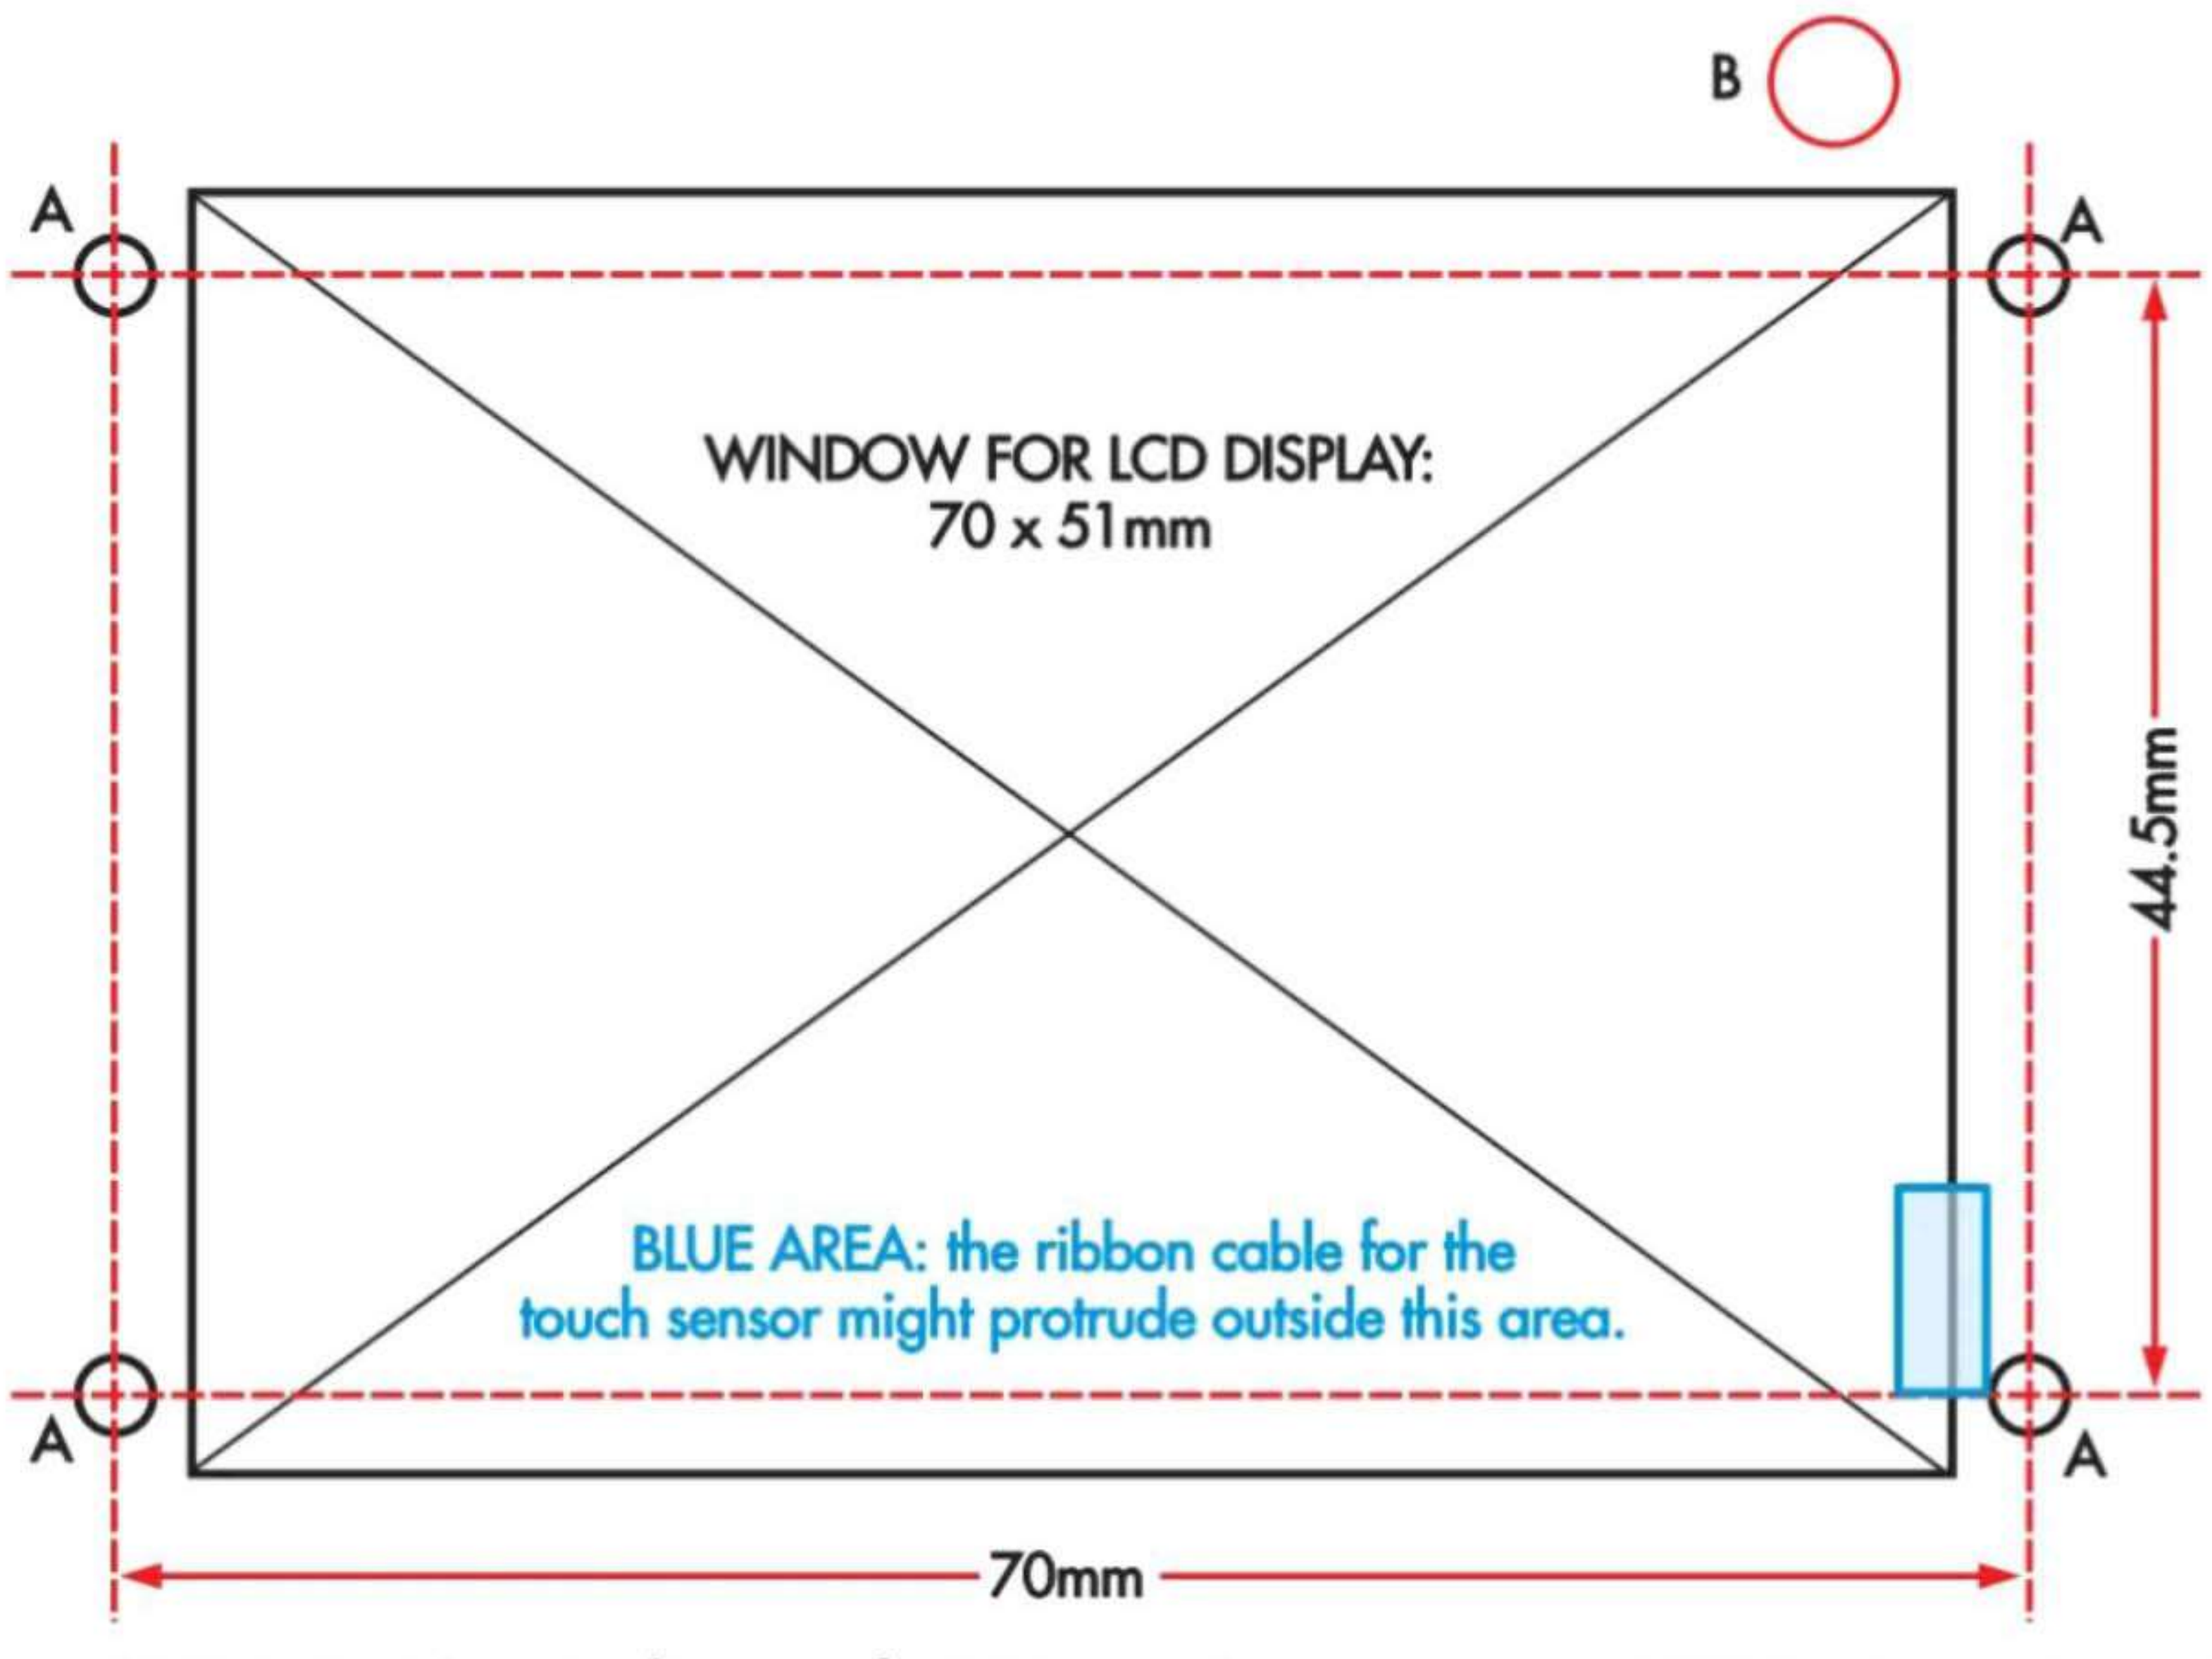
B
A
A
WINDOW FOR LCD DISPLAY:
70 x 51mm
44.5mm
BLUE AREA: the ribbon cable for the touch sensor might protrude outside this area.
A
A
70mm

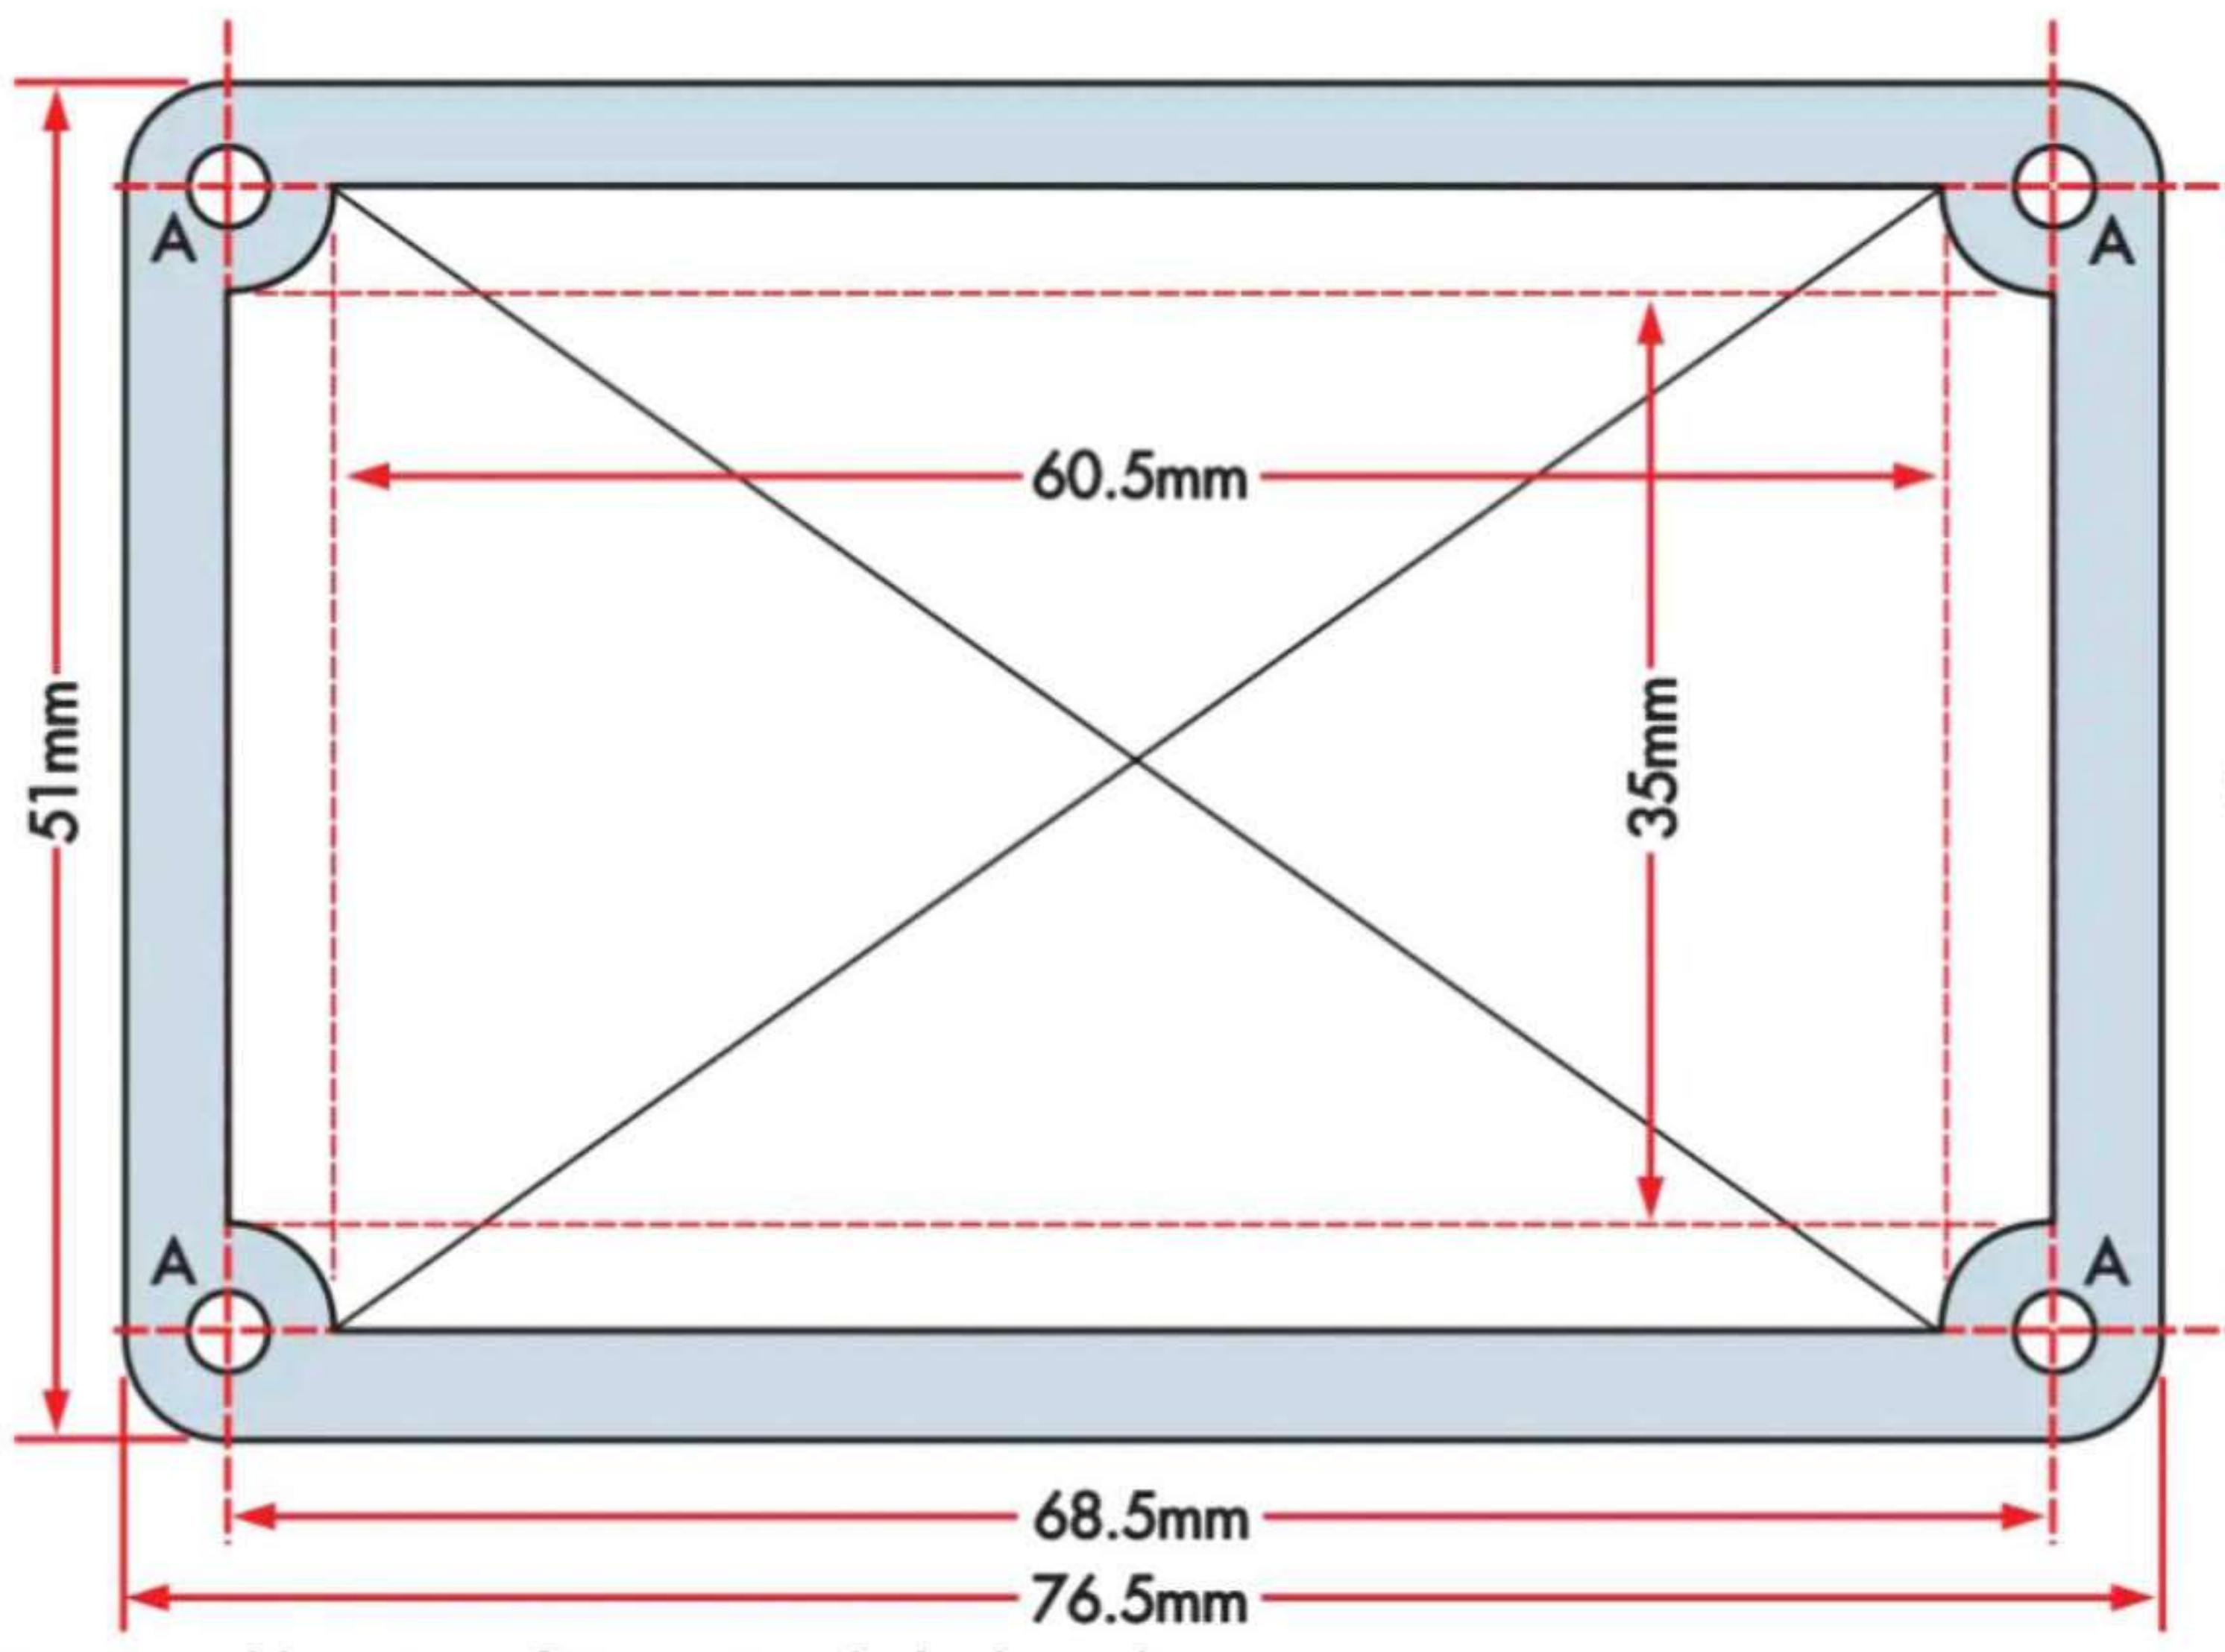

HOLES A: 3.0mm in diameter, for LCD mounting screws

HOLE B: Approximate size (5.0mm) and location of IR receiver hole, based on our prototype

Fig.12 (left): to mount the 2.8in LCD and thus the Control Module, you'll need a large rectangular hole and four small round holes. You might also need another small hole for the IR receiver to 'see' outside (like the one marked "B").

Fig.13 (right): the exact dimensions of the cutout for the OLED Module are not critical, as the shape overlaps the edge of the hole by about 4mm. Still, you might need to use foam tape or similar to protect the back of the PCB if you are using a metal enclosure.

DOWNLOAD YOUR PE/EPE BACK ISSUES

COMPLETE-YEAR collections – all 12 issues for each year!

From just £5.95 each

VOL 28 – January to December 1999 – **£5.95**
VOL 29 – January to December 2000 – **£6.95**
VOL 30 – January to December 2001 – **£7.95**
VOL 31 – January to December 2002 – **£7.95**
VOL 32 – January to December 2003 – **£7.95**
VOL 33 – January to December 2004 – **£7.95**
VOL 34 – January to December 2005 – **£8.95**
VOL 35 – January to December 2006 – **£8.95**
VOL 36 – January to December 2007 – **£8.95**
VOL 37 – January to December 2008 – **£8.95**
VOL 38 – January to December 2009 – **£10.95**
VOL 39 – January to December 2010 – **£11.95**
VOL 40 – January to December 2011 – **£12.95**
VOL 41 – January to December 2012 – **£14.95**
VOL 42 – January to December 2013 – **£15.95**
VOL 43 – January to December 2014 – **£16.95**
VOL 44 – January to December 2015 – **£17.95**
VOL 45 – January to December 2016 – **£19.95**
VOL 46 – January to December 2017 – **£22.95**
VOL 47 – January to December 2018 – **£27.95**
VOL 48 – January to December 2019 – **£27.95**
VOL 49 – January to December 2020 – **£29.95**
VOL 50 – January to December 2021 – **£32.95**
VOL 51 – January to December 2022 – **£34.95 – NEW!**

FIVE-YEAR collections – 60 issues!

From just £14.95 each

PE2832: Jan 99 to Dec 03 – £14.95
PE2933: Jan 00 to Dec 04 – £14.95
PE3034: Jan 01 to Dec 05 – £14.95
PE3135: Jan 02 to Dec 06 – £14.95
PE3236: Jan 03 to Dec 07 – £14.95
PE3337: Jan 04 to Dec 08 – £14.95
PE3438: Jan 05 to Dec 09 – £19.95
PE3539: Jan 06 to Dec 10 – £25.95
PE3640: Jan 07 to Dec 11 – £25.95
PE3741: Jan 08 to Dec 12 – £25.95
PE3842: Jan 09 to Dec 13 – £33.95
PE3943: Jan 10 to Dec 14 – £33.95
PE4044: Jan 11 to Dec 15 – £33.95
PE4145: Jan 12 to Dec 16 – £35.95
PE4246: Jan 13 to Dec 17 – £37.95
PE4347: Jan 14 to Dec 18 – £38.95
PE4448: Jan 15 to Dec 19 – £39.95
PE4549: Jan 16 to Dec 20 – £42.95
PE4650: Jan 17 to Dec 21 – £44.95 – NEW!

20 YEARS OF PE!

AMAZING BUNDLE OFFER
Every issue from 1999 – 2018

240 magazines for just

£59.95

PE2832: Jan 99 to Dec 03
PLUS!
PE3337: Jan 04 to Dec 08
PLUS!
PE3842: Jan 09 to Dec 13
PLUS!
PE3447: Jan 14 to Dec 18

Do you have the 15-year bundle?

Upgrade to the new 20-year bundle for just £14.95 with: PE15to20

Download your collections at: www.electronpublishing.com

Practical Electronics PCB SERVICE

PROJECT	CODE	PRICE
DECEMBER 2024		
Raspberry Pi Clock Radio PCB pair	19101241/2	£12.99
VGA PicoMite (blue)	07107221	£5.99
Secure Remote Switch receiver	10109231	£4.99
Secure Remote Switch transmitter (module version)	10109232	£2.99
Secure Remote Switch transmitter (discrete version)	10109233	£2.99
NOVEMBER 2024		
Ideal Diode Bridge Rectifiers (28mm square)	18101241	£3.99
Ideal Diode Bridge Rectifiers (22mm square)	18101242	£3.99
Ideal Diode Bridge Rectifiers (4-pin SIL)	18101243	£3.99
Ideal Diode Bridge Rectifiers (tiny SOT-23)	18101244	£3.99
Ideal Diode Bridge Rectifiers (standalone SMD)	18101245	£4.99
Ideal Diode Bridge Rectifiers (standalone TH; 2oz copper)	18101246	£6.99
Multi-channel Volume Control (volume PCB)	01111221	£4.99
Multi-channel Volume Control (control PCB)	01111222	£4.99
Multi-channel Volume Control (OLED PCB, black)	01111223	£4.99
OCTOBER 2024		
TQFP-32 Programming Adaptor	24108231	£4.99
TQFP-44 Programming Adaptor	24108232	£4.99
TQFP-48 Programming Adaptor	24108233	£4.99
TQFP-64 Programming Adaptor	24108234	£4.99
2m VHF FM Test Signal Generator	06107231	£4.99
SEPTEMBER 2024		
Calibrated Microphone (SMD version)	01108231	£3.99
Calibrated Microphone (TH version)	01108232	£3.99
PIC Programming Adaptor	24105231	£4.99
Salad Bowl Speaker crossover	01109231	£7.99
AUGUST 2024		
Arduino ESR Meter combined L/C-ESR	04106182	£9.99
Arduino ESR Meter separate ESR	04106183	£7.99
Watering System Controller	15110231	£11.99
JULY 2024		
Dynamic NFC Tag (22 × 31mm)	06101231	£5.95
Dynamic NFC Tag (86 × 54mm)	06101232	£7.95
Reciprocal Frequency Counter	CSE230101C	£8.95
Reciprocal Frequency Counter front panel (black mask)	CSE230102	£7.95
Pi Pico-based Thermal Camera	04105231	£6.95
JUNE 2024		
Loudspeaker Test Jig	04106231	£12.95
MAY 2024		
GPS-Disciplined Oscillator	04103231	£12.95
Dual RF Amplifier	CSE220602A	£5.95
Songbird	08103231	£7.95
APRIL 2024		
Basic RF Signal Generator	CSE221001	£8.99
Basic RF Signal Generator front panel (black solder mask)	CSE220902B	£9.99
Silicon Chirp Cricket	08101231	£7.99
Test Bench 'Swiss Army Knife'	04110221	£9.99
MARCH 2024		
Digital Volume Control Potentiometer SMD version	01101231	£8.95
Digital Volume Control Potentiometer through-hole version	01101232	£9.95
Advanced SMD Test Tweezers body	04106221	£6.95
Advanced SMD Test Tweezers legs (pair)	04106212	£7.95
FEBRUARY 2024		
Active Mains Soft Starter	10110221	£10.95

PROJECT	CODE	PRICE
JANUARY 2024		
Q Meter	CSE220701	£8.95
Q Meter front panel (black solder mask)	CSE220704	£7.95
Raspberry Pi Pico W BackPack	07101221	£8.95
DECEMBER 2023		
Digital Boost Regulator	24110224	£8.95
Dual-Channel Power Supply for Breadboards	04112221	£11.95
Display Adaptor for the Breadboard PSU	04112222	£9.95
NOVEMBER 2023		
LC Meter Mk3	CSE220503C	£9.95
DC Supply Filter for vehicles	08108221	£8.95
Discrete Audio Op Amp PCB	AO1-JUL23	£9.95
Discrete Audio Op Amp PCB + essential components	AO2-JUL23	£17.95
OCTOBER 2023		
Buck/Boost Charger Adaptor	14108221	£9.95
PIC Breakout Board for SOIC parts	24110225	£5.95
PIC Breakout Board for DIP parts	24110222	£5.95
AVR64DD32 Breakout board	24110223	£5.95
Automatic Train Controller	09109221	£6.95
Chuff Sound module	09109222	£6.95
SEPTEMBER 2023		
Mini LED Driver	16106221	£7.95
New GPS-Synchronised Clock	19109221	£9.95
AUGUST 2023		
110dB RF Attenuator	CSE211003	£9.95
Universal Battery Charge Controller (2023 update)	14107192	£9.95
Wide-Range OhmMeter	04109221	£12.95
JULY 2023		
Multimeter Checker/Calibrator	04107221	£10.95
JUNE 2023		
Arduino Programmable Load	04105221	£8.95
Buck-Boost LED Driver	16103221	£8.95
MAY 2023		
Precision AM-FM DDS Signal Generator	CSE211002	£14.95
Improved SMD Test Tweezers programmed PIC	0410621PIC	£12.95
AO 2x-dual-to-quad through-hole (pack of five)	AO1-MAY23	£6.95
AO 2x-dual-to-quad SMD (pack of five)	AO2-MAY23	£6.95
AO 2x-single-to-dual through-hole (pack of five)	AO3-MAY23	£6.95
AO 2x-single-to-dual SMD (pack of five)	AO4-MAY23	£6.95
APRIL 2023		
500W Amplifier Module	see p.22, April 2023	
Clipping Indicator (per channel)	01112211	£7.95
CD Welder Power Supply (one needed)	29103221	£9.95
CD Welder Controller (one needed)	29103222	£9.95
CD Welder Energy Storage module (several needed)	29103223	£7.95
AO Universal Dual Op Amp Board	AO1-APR23	£9.95
AO Stereo RIAA precision passives kit for dual op amp	AO2-APR23	£8.95
MARCH 2023		
Pico BackPack	07101221	£9.95
Semaphore Signal (controller)	09103221	£7.95
Semaphore Signal (blade)	09103222	£5.95
FEBRUARY 2023		
Tesla Coil driver board	26102221	£9.95
Tesla Coil potentiometer board	26102222	£5.95
Cooling Fan Controller & Loudspeaker Protector	01102221	£8.95
Remote Gate Controller	11009121	£12.95

PCBs for most recent *PE/EPE* constructional projects are available. From the July 2013 issue onwards, PCBs with eight-digit codes have silk screen overlays and, where applicable, are double-sided, have plated-through holes, and solder mask. They are similar to photos in the project articles. Earlier PCBs are likely to be more basic and may not include silk screen overlay, be single-sided, lack plated-through holes and solder mask.

Always check price and availability in the latest issue or online. A large number of older boards are listed for ordering on our website.

In most cases we do not supply kits or components for our projects. For older projects it is important to check the availability of all components before purchasing PCBs.

Back issues of articles are available – see Back Issues page for details.

Double-sided | plated-through holes | solder mask

PROJECT	CODE	PRICE
JANUARY 2023		
Classic LED Metronome – 8-LED	23111211	£7.95
Classic LED Metronome – 10-LED	23111212	£8.95
Multi-Channel Speaker Protector – 6 channel	01101221	£9.95
Multi-Channel Speaker Protector – 4 channel	01101222	£7.95
Remote Control Range Extender – IR-to-UHF	15109212	£5.95
Remote Control Range Extender – UHF-to-IR	15109211	£7.95
AO Universal Single Op Amp Board	AO1-JAN23	£7.95
DECEMBER 2022		
Hummingbird Amplifier	01111211	£9.95
SMD Trainer PCB	29106211A	£8.95
SMD Trainer PCB + parts	29106211B	£13.95
NOVEMBER 2022		
USB Cable tester – main PCB	04108211	£12.95
USB Cable tester – front panel	04108212	£5.95
USB Cable tester – optional panel	SC5970	£5.95
Model Railway Carriage Lights – PCB	09109211	£6.95
AO trnsfmr PCB – standard VTX-A range	VTX-101-007	£6.95
AO trnsfmr PCB – dual-outline VTX102-3001/101-3002	VTX-Dual	£6.95
OCTOBER 2022		
SMD Test Tweezers – PCB and pair of tweezer arms	04106211-2	£11.95
SMD Test Tweezers – programmed PIC12F1572-I/SN	0410621A	£7.95
Tele-com	12110211	£12.95
SEPTEMBER 2022		
Touchscreen Digital Preamp – main board	01103191	£12.95
Touchscreen Digital Preamp – adaptor board pair	01103192	£5.95
20A DC Motor Speed Controller	11006211	£9.95
AUGUST 2022		
Multi-purpose Battery Manager – I/O Expander	11104212	£5.95
Multi-purpose Battery Manager – Switch Module	11104211	£8.95
Simple MIDI Music Keyboard (for 8 switches)	23101213	£6.95
Nano Pong	08105212	£7.95
JULY 2022		
Silicon Labs AM/FM/SW Radio	CSE210301C	£10.95
Level Crossing Controller	09108211	£6.95
JUNE 2022		
Full-wave Motor Speed Controller	1010221	£8.95
PIC Programming Helper for 8-pin PICs only	24106211	£7.95
PIC Programming Helper for 8, 14 or 20-pin PICs	24106212	£10.95
Advanced GPS Computer	05102211	£9.95
MAY 2022		
Bus board PCB for Analogue Vocoder	AO1-MAY22	£10.95
Complete set of 14 PCBs for Analogue Vocoder	AO2-MAY22	£97.95
Programmed EEPROM for Digital FX Unit	FX1-MAY22	£10.95
Programmed PIC for Digital FX Unit w/ potentiometer	FX2-MAY22	£8.95
APRIL 2022		
64-key MIDI Matrix shield	23101211	£8.95
64-key MIDI Matrix switch board	23101212	£11.95
High-current Battery Balancer	14102211	£10.95
Digital FX Unit – using potentiometer	01102211	£9.95
Digital FX Unit – using BCD switch	01102212	£9.95
Universal Audio PSU	AO1-APR22	£11.95
MARCH 2022		
Mini Isolated Serial Link	24102211	£5.95
Busy Loo Indicator	16112201	£5.95
Analogue Vocoder – Band-pass filter board	AO1-MAR22	£9.95
Analogue Vocoder – HP/LP filter board	AO2-MAR22	£9.95

PROJECT	CODE	PRICE
FEBRUARY 2022		
Arduino-based Power Supply	18106201	£9.95
Battery Monitor Logger	11106201	£10.95
Electronic Wind Chimes	23011201	£10.95
Analogue Vocoder – Driver Amplifier	AO-FEB22	£8.95
JANUARY 2022		
Vintage battery Radio Li-ion Power Supply	11111201	£9.95
MiniHeart: A Miniature Heartbeat Simulator	01109201	£8.95
DECEMBER 2021		
AM/FM/SW Digital Receiver	CSE200902A	£13.95
Balanced Input and Attenuator for USB CODEC	01106202	£11.95
NOVEMBER 2021		
Dual Battery Lifesaver	11111202	£6.95
OCTOBER 2021		
Mini Wi-Fi LCD BackPack	24106201	£8.95

For the many earlier PCBs that we stock please see the PE website: www.electronpublishing.com

PE/EPE PCB SERVICE

Order Code	Project	Quantity	Price

Name ....

Address ....

....

Tel ....

Email ....

Card No ....

Valid From ....Expiry Date ....

Card Security No ....

You can also order PCBs by phone, email or via the shop on our website: www.electronpublishing.com

No need to cut your issue – a copy of this form is just as good!

All prices include VAT and UK p&p. Add £4 per *project* for post to Europe; £5 per *project* outside Europe.

Orders and payment should be sent to:

Practical Electronics, Electron Publishing Ltd
113 Lynwood Drive, Merley, Wimborne, Dorset BH21 1UU
Tel 01202 880299 Email: pe@pemag.au
On-line Shop: www.electronpublishing.com

Sorry, cheques cannot be accepted as we are based in Australia and our UK bank account does not allow cheque deposits.

NOTE: Most boards are in stock and sent within seven days of order receipt. Please allow up to 28 days delivery if we need to restock.

Practical Electronics PCB SERVICE

PROJECT	CODE	PRICE
SEPTEMBER 2021		
USB SuperCodec PCB	01106201	£14.95
Audio DDS Oscillator PCB	01110201	£5.95
Audio DDS Oscillator rotary encoder	01110201-ENC	6.95
Programming Adaptor Board for Audio DDS Oscillator	01110202	£5.95
High-power Ultrasonic Cleaner main PCB	04105201	£14.95
High-power Ultrasonic Cleaner front-panel PCB	04105202	
Night Keeper Lighthouse PCB	08110201	£5.95
AUGUST 2021		
Ol' Timer PCB	19104201	£11.95
Ol' Timer 8x8 RGB LED module using WS2812B	19104201-88	£8.95
Ol' Timer set of acrylic case pieces and spacer	19104201-ACR	£8.75
Ol' Timer DS3231 RTC IC wide SOIC-16	19104201-RTC	£5.95
Wideband Digital RF Power Meter	04106201	£9.75
Switchmode 78xx regulators ***(PACK of 5!)***	18105201	£7.95
Cool Beans SMAD display	CB-AUG21	£11.95
JULY 2021		
ATtiny816 Breakout / Dev Board w/ Capacitive Touch	24110181	£9.75
IR Remote Control Assistant (Jaycar version)	15005201	£8.95
IR Remote Control Assistant (Altronics version)	15005202	£8.95
PIC18F Development Board	PNM-JUL21	£12.95
Microphone Preamplifier	AO-JUL21	£11.95
JUNE 2021		
Roadies' Test Signal Generator (surface-mount ver.)	01005201	£8.95
Roadies' Test Signal Generator (through-hole version)	01005202	£9.95
Touchscreen Wide-range RCL Box (Resistor module)	04104201	£18.95
Touchscreen Wide-range RCL Box (Ind/Cap module)	04104202	
KickStart Part 3 – Gyrator-based Audio Filter	KS3-2021	£7.95
MAY 2021		
7-Band Equaliser (Mono)	01104201	£8.95
7-Band Equaliser (Stereo)	01104202	£10.95
Car Altimeter	05105201	£7.95
APRIL 2021		
Reflow Oven – DSP Active Crossover (CPU)	01106193	£19.95
Reflow Oven – DSP Active Crossover (Front panel)	01106195	
Reflow Oven – DSP Active Crossover (LCD)	01106196	
Frequency Reference Signal Distributor	CSE200103	£8.95
MARCH 2021		
Nutube Guitar Effects Pedal	01102201	£12.95
Programmable Thermal Regulator (Peltier Interface)	21109181	£18.95
Programmable Thermal Regulator (Peltier Driver)	21109182	
Tunable HF Preamp	CSE190502	£8.95
FEBRUARY 2021		
4G Remote Monitoring	27111191	£9.95
JANUARY 2021		
Nutube Valve Preamplifier	01112191	£12.95
Arduino DCC Controller	09207181	£10.95
DECEMBER 2020		
Pseudo-Random Sequence Generator	16106191	£7.95
Clever Charger	14107191	£11.95
PE Theremin Amplifier	AO-1220-01	£8.95

PROJECT	CODE	PRICE
NOVEMBER 2020		
LED Christmas Tree (1 off)	16107181-1	£6.95
LED Christmas Tree (4 off)	16107181-2	£14.95
LED Christmas Tree (12 off)	16107181-3	£24.95
LED Christmas Tree (20 off)	16107181-4	£34.95
USB/SPI Interface Board	16107182	£5.95
45V/8A Power Supply PCB *plus* acrylic spacer	18111181	£14.95
45V/8A Power Supply five-way display bezel	18111181-BZ	£3.95
Five-way LCD Panel Meter/Display	18111182	£7.95
OCTOBER 2020		
Digital Audio Millivoltmeter	04108191	£9.95
Precision Signal Amplifier	04107191	£7.95
SEPTEMBER 2020		
PE Theremin PSU	AO-0920-01	£5.95
PE Theremin PSU transformer	AO-0920-02	£7.95
Micromite Explore-28	07108191	£6.95
Ultrabrite LED Driver	16109191	£6.95
AUGUST 2020		
Micromite LCD BackPack V3	07106191	£9.95
Steering Wheel Audio Button to Infrared Adaptor	05105191	£7.95
JULY 2020		
AM/FM/CW Scanning HF/VHF RF Signal Generator	04106191	£13.95
Speech Synthesiser with the Raspberry Pi Zero	01106191	£5.95
PE Mini-organ PCB	AO-0720-01	£14.95
PE Mini-organ selected parts	AO-0720-02	£8.95
High-current 12V Battery Isolator – control	05106191	£6.95
High-current 12V Battery Isolator FET (2oz)	05106192	£9.95
JUNE 2020		
Arduino breakout board – 3.5-inch LCD Display	24111181	£6.95
Six-input Audio Selector main board	01110191	£10.95
Six-input Audio Selector switch panel board	01110192	
MAY 2020		
Ultra-low-distortion Preamplifier Input Selector	01111112	£11.25
Ultra-low-distortion Preamplifier pushbutton board	01111113	
Universal Regulator	18103111	£7.95
433MHz Wireless Data Repeater	15004191	£8.50
Bridge-mode Adaptor for Amplifier	01105191	£7.95
iCEstick VGA Terminal	02103191	£5.95
Analogue noise with tilt control	AO-0520-01	£7.95
Audio Spectrum Analyser	PM-0520-01	£8.95
APRIL 2020		
Flip-dot Display black coil board	19111181	£14.95
Flip-dot Display black pixels	19111182	
Flip-dot Display black frame	19111183	
Flip-dot Display green driver board	19111184	
MARCH 2020		
Diode Curve Plotter	04112181	£10.95
Steam Train Whistle / Diesel Horn Sound Generator	09106181	£8.50
Universal Passive Crossover (one off)	UPC0320	£12.50
FEBRUARY 2020		
Motion-Sensing 12V Power Switch	05102191	£5.95
USB Keyboard / Mouse Adaptor	24311181	£8.50
DSP Active Crossover (ADC)	01106191	£29.95
DSP Active Crossover (DAC) ×2	01106192	
DSP Active Crossover (CPU)	01106193	
DSP Active Crossover (Power/routing)	01106194	
DSP Active Crossover (Front panel)	01106195	
DSP Active Crossover (LCD)	01106196	

Silicon Chip kits & parts for projects in this issue (posted from Australia):

VGA PicoMite short-form kit without case (SC6417)	https://siliconchip.au/Shop/20/6417	AUD $37.50 (about £19) + P&P
VGA PicoMite short-form kit with case (SC7402)	https://siliconchip.au/Shop/20/7402	AUD $55.00 (about £28) + P&P
Secure Remote Switch receiver kit (SC7401)	https://siliconchip.au/Shop/20/7401	AUD $75.00 (about £39) + P&P
Secure Remote Switch discrete transmitter kit (SC6836)	https://siliconchip.au/Shop/20/6836	AUD $20.00 (about £10) + P&P
Secure Remote Switch module transmitter kit (SC6837)	https://siliconchip.au/Shop/20/6837	AUD $15.00 (about £8) + P&P
Multi-channel Volume Control control module kit (SC6793)	https://siliconchip.au/Shop/20/6793	AUD $50.00 (about £26) + P&P
Multi-channel Volume Control volume module kit (SC6794)	https://siliconchip.au/Shop/20/6794	AUD $55.00 (about £28) + P&P
Multi-channel Volume Control OLED module kit (SC6795)	https://siliconchip.au/Shop/20/6795	AUD $25.00 (about £13) + P&P

20% VAT is added for customers in the UK. Postage & packing cost is typically around £12-15.
Please allow 1-2 weeks for delivery (longer in some cases).

CLASSIFIED ADVERTISING

Practical Electronics

If you want your advertisements to be seen by the largest readership at the most economical price then our classified page offers excellent value.

The rate for semi-display space is £10 per centimetre high, with a minimum height of 2·5cm. All semi-display advertisements have a width of 5.5cm. The prepaid rate for classified adverts is 40p per word (minimum 12 words).

For rates, further information on display and classified advertising or to book an ad, please email us at **pe@pemag.au**

Practical Electronics reaches more UK readers than any other UK monthly hobby electronics magazine. Our sales figures prove it.

We have been the leading monthly magazine in this market for the last twenty-seven years.

BOWOOD ELECTRONICS LTD
Suppliers of Electronic Components
www.bowood-electronics.co.uk
Unit 10, Boythorpe Business Park, Dock Walk, Chesterfield,
Derbyshire S40 2QR. Sales: 01246 200 222
Send large letter stamp for Catalogue

Electrical Industries Charity (EIC)
We help people working in the electrical, electronics and energy community as well as their family members and retirees.

We use workplace programmes that give the industry access to financial grants and a comprehensive range of free and confidential services.
www.electricalcharity.org

electrical industries charity

COAST ELECTRONICS
BREAKOUTS-COMPONENTS-CONTRACT DESIGN-3D PRINTER PARTS-MUSICAL-MICROCONTROLLERS
WWW.COASTELECTRONICS.CO.UK

MISCELLANEOUS

VALVES AND ALLIED COMPONENTS?
For free stock list and/or advice, please contact me: geoffdavies337@gmail.com Telephone: 01788 574774

PIC DEVELOPMENT KITS, DTMF kits and modules, CTCSS Encoder and Decoder/Display kits.
Visit www.cstech.co.uk

COMPONENTS AS JOB LOT
Large stock. Transistors, Diodes, Logic, Memory and Analogue parts. 1000s of passive items. Catalogued and sorted. 'New Old Stock' in racks.
Email: rojobarn@gmail.com

Andrew Kenny – Qualified Patent Agent
EPO UKIPO USPTO
Circuits Electric Machinery Mechatronics
Web: www.akennypatentm.com
Email: Enquiries@akennypatentm.com
Tel: 0789 606 9725

ADVERTISING INDEX

ENIGMA COMPONENTS SHOP 35
FLOWCODE .... Inner front cover
HAMMOND ELECTRONICS Ltd. 6
MICROCHIP TECHNOLOGY .... Outer back cover
PCBWAY. 7
STEWART OF READING 53

Advertising mailing address
Electron Publishing Australia
PO Box 194
Matraville NSW 2036
Australia
Email pe@pemag.au
Web www.electronpublishing.com

For editorial contact details, see page 2.

Next Month – in the January issue

High-performance Microphone Preamplifier

9V DC powered for stage and studio work. It fits into a sturdy diecast aluminium box or can be integrated into other equipment.

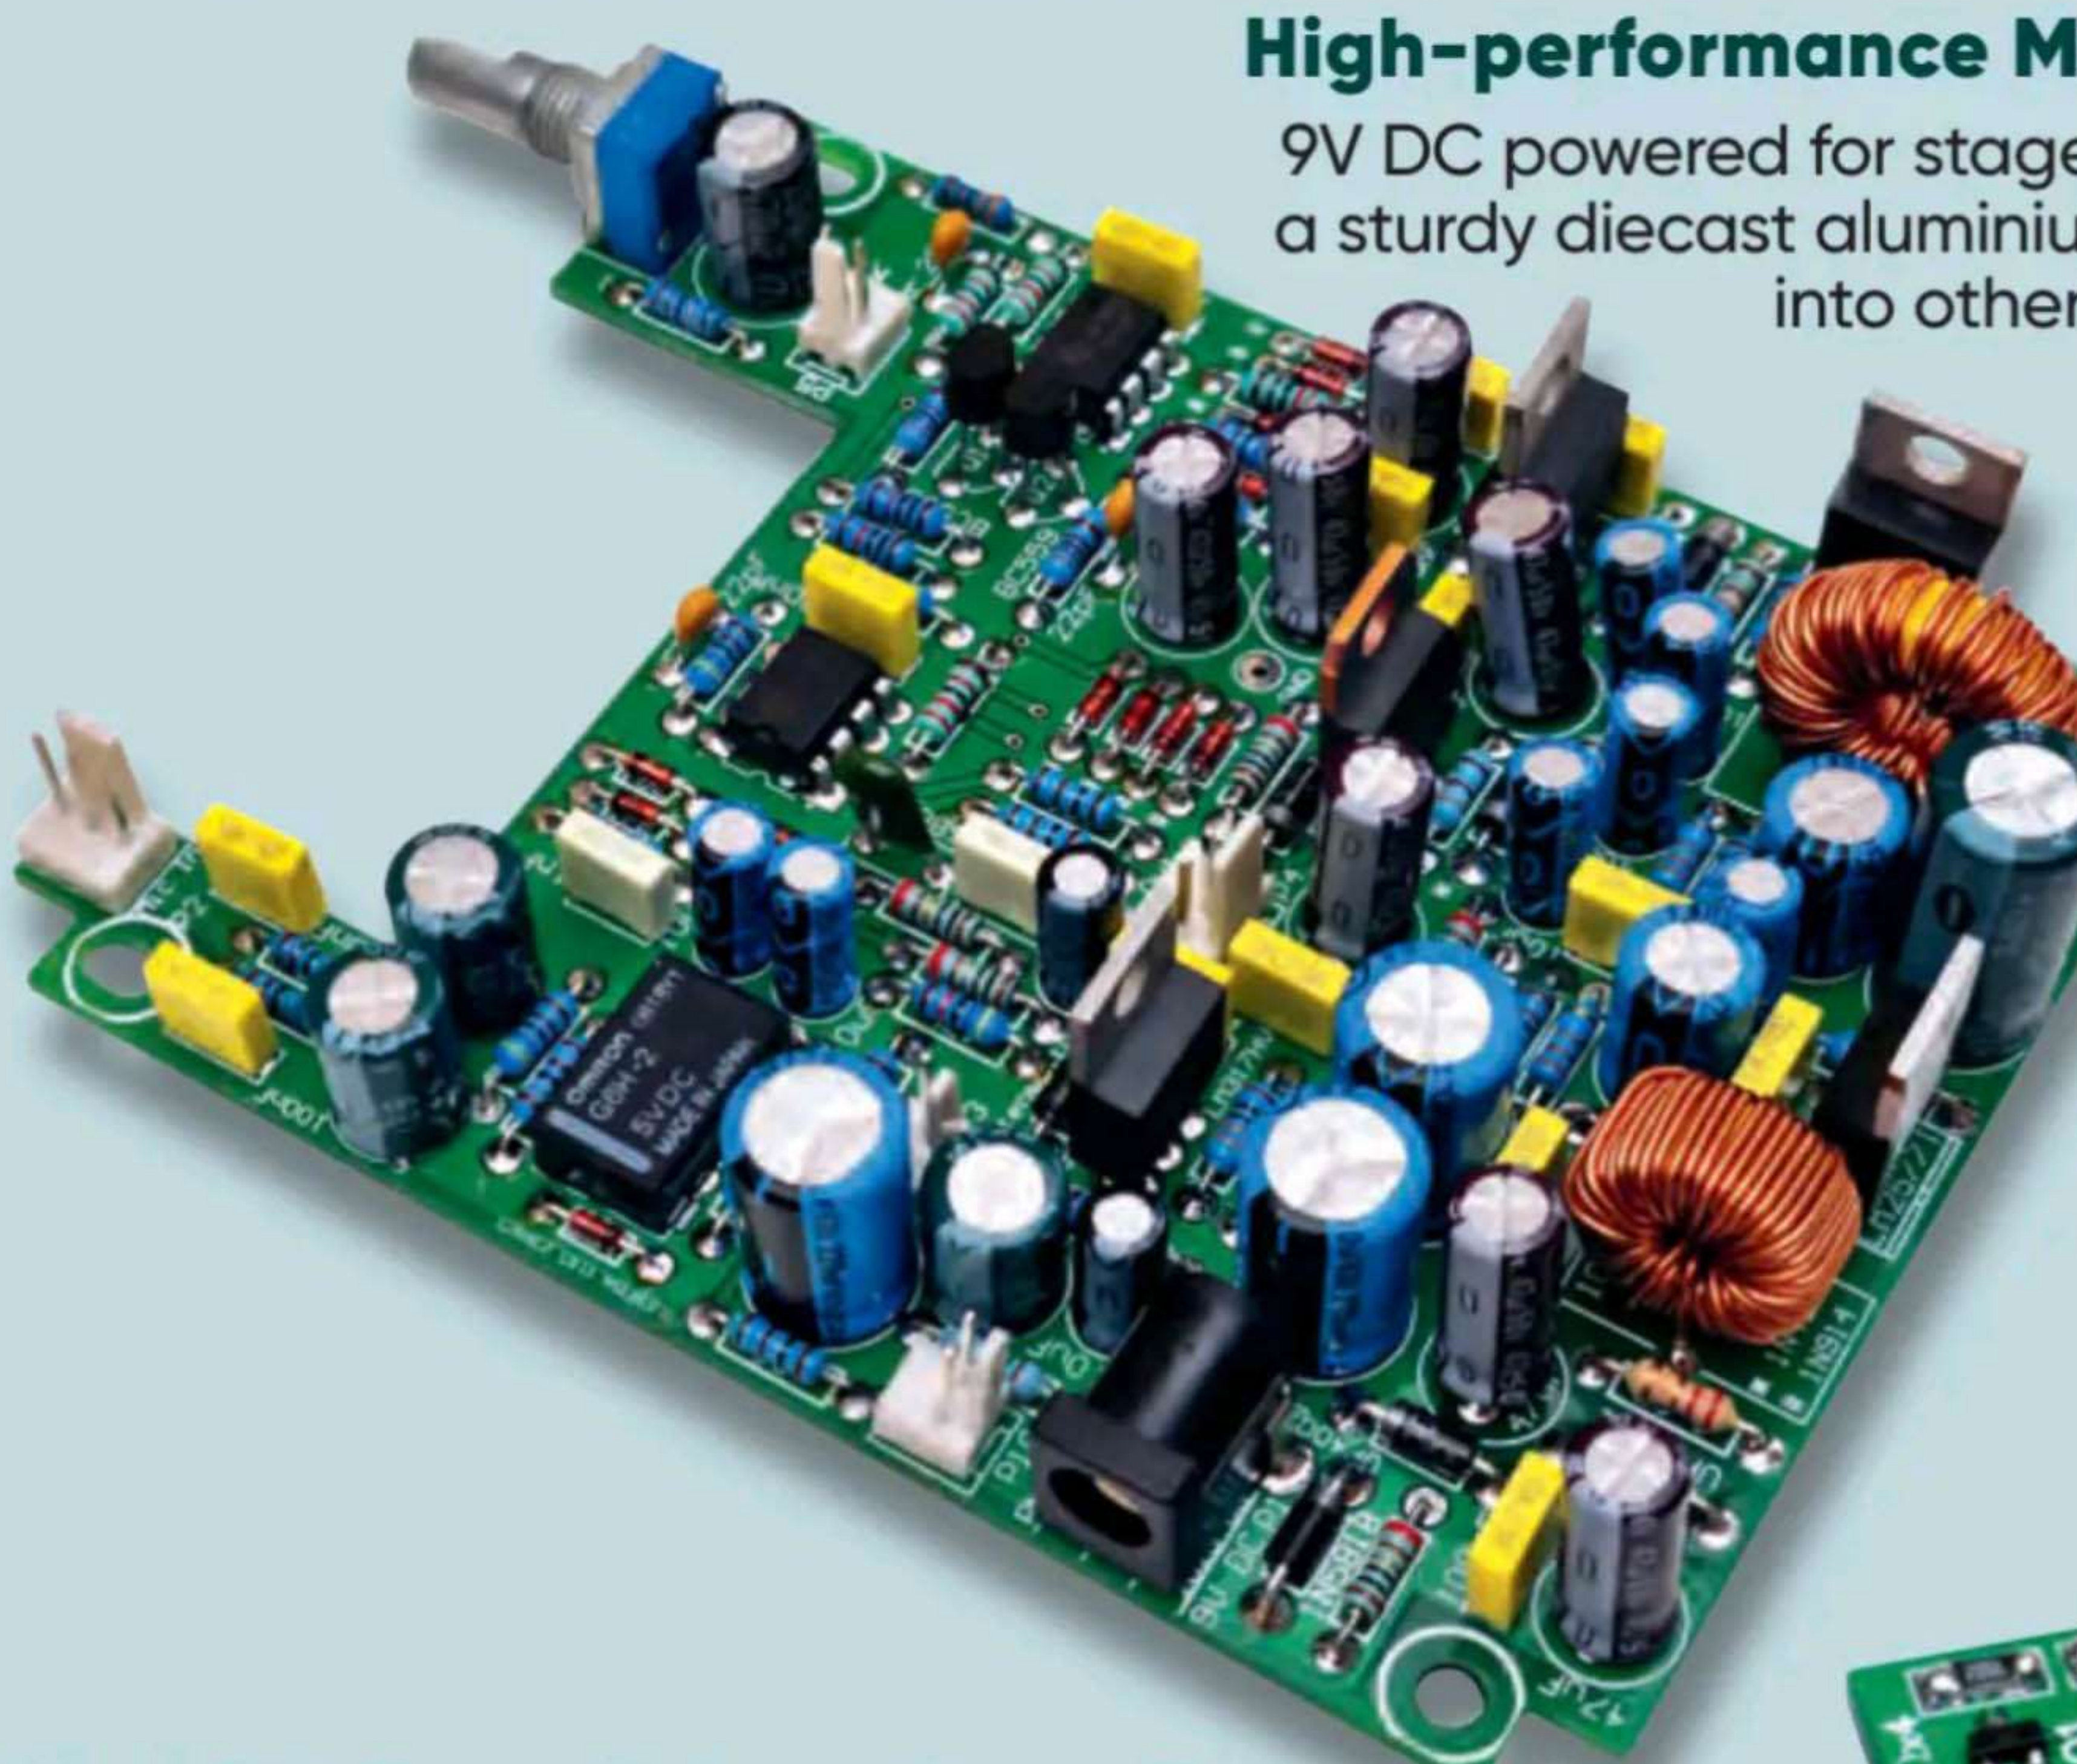

Model Railway Points Controller

Up to eight sets per controller PCB, with optional LED-based signals.

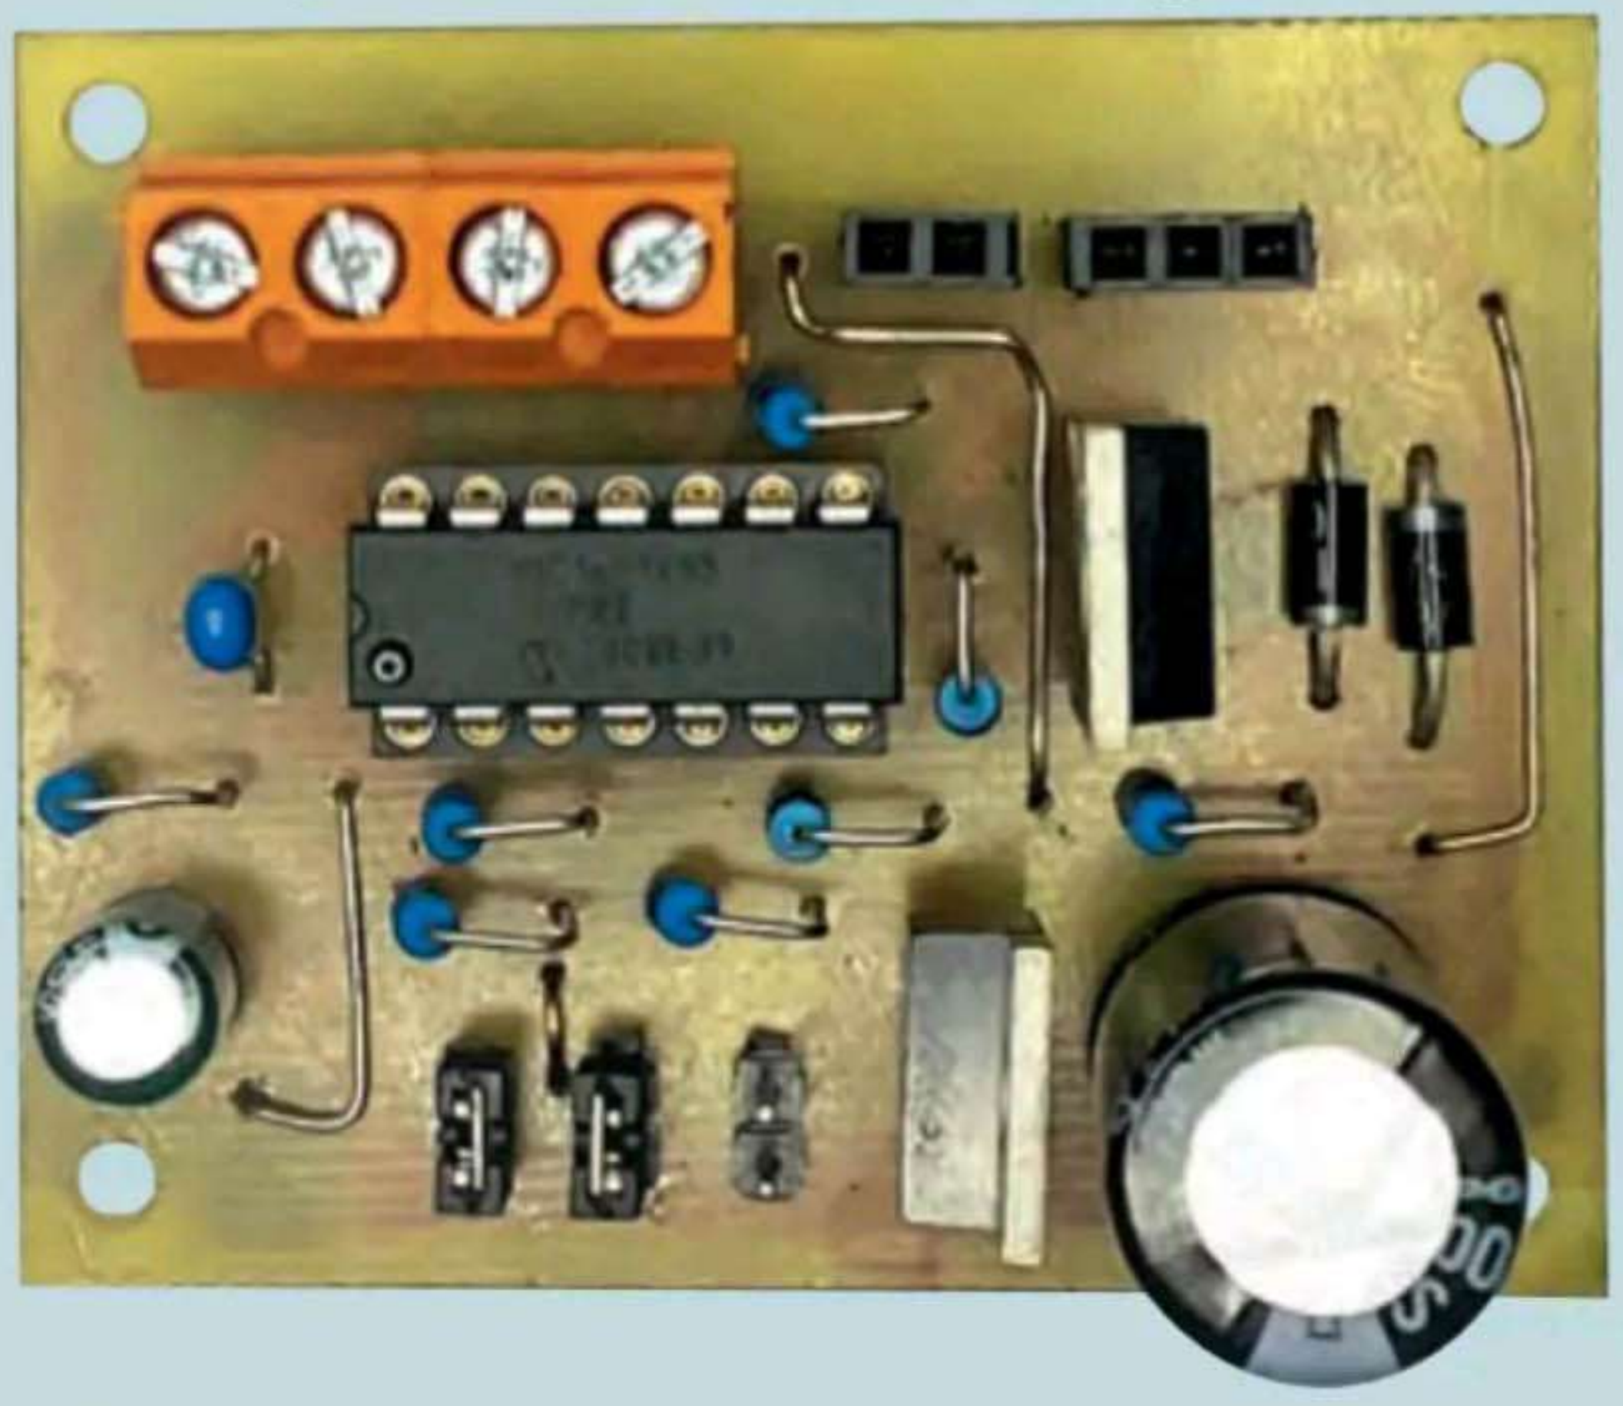

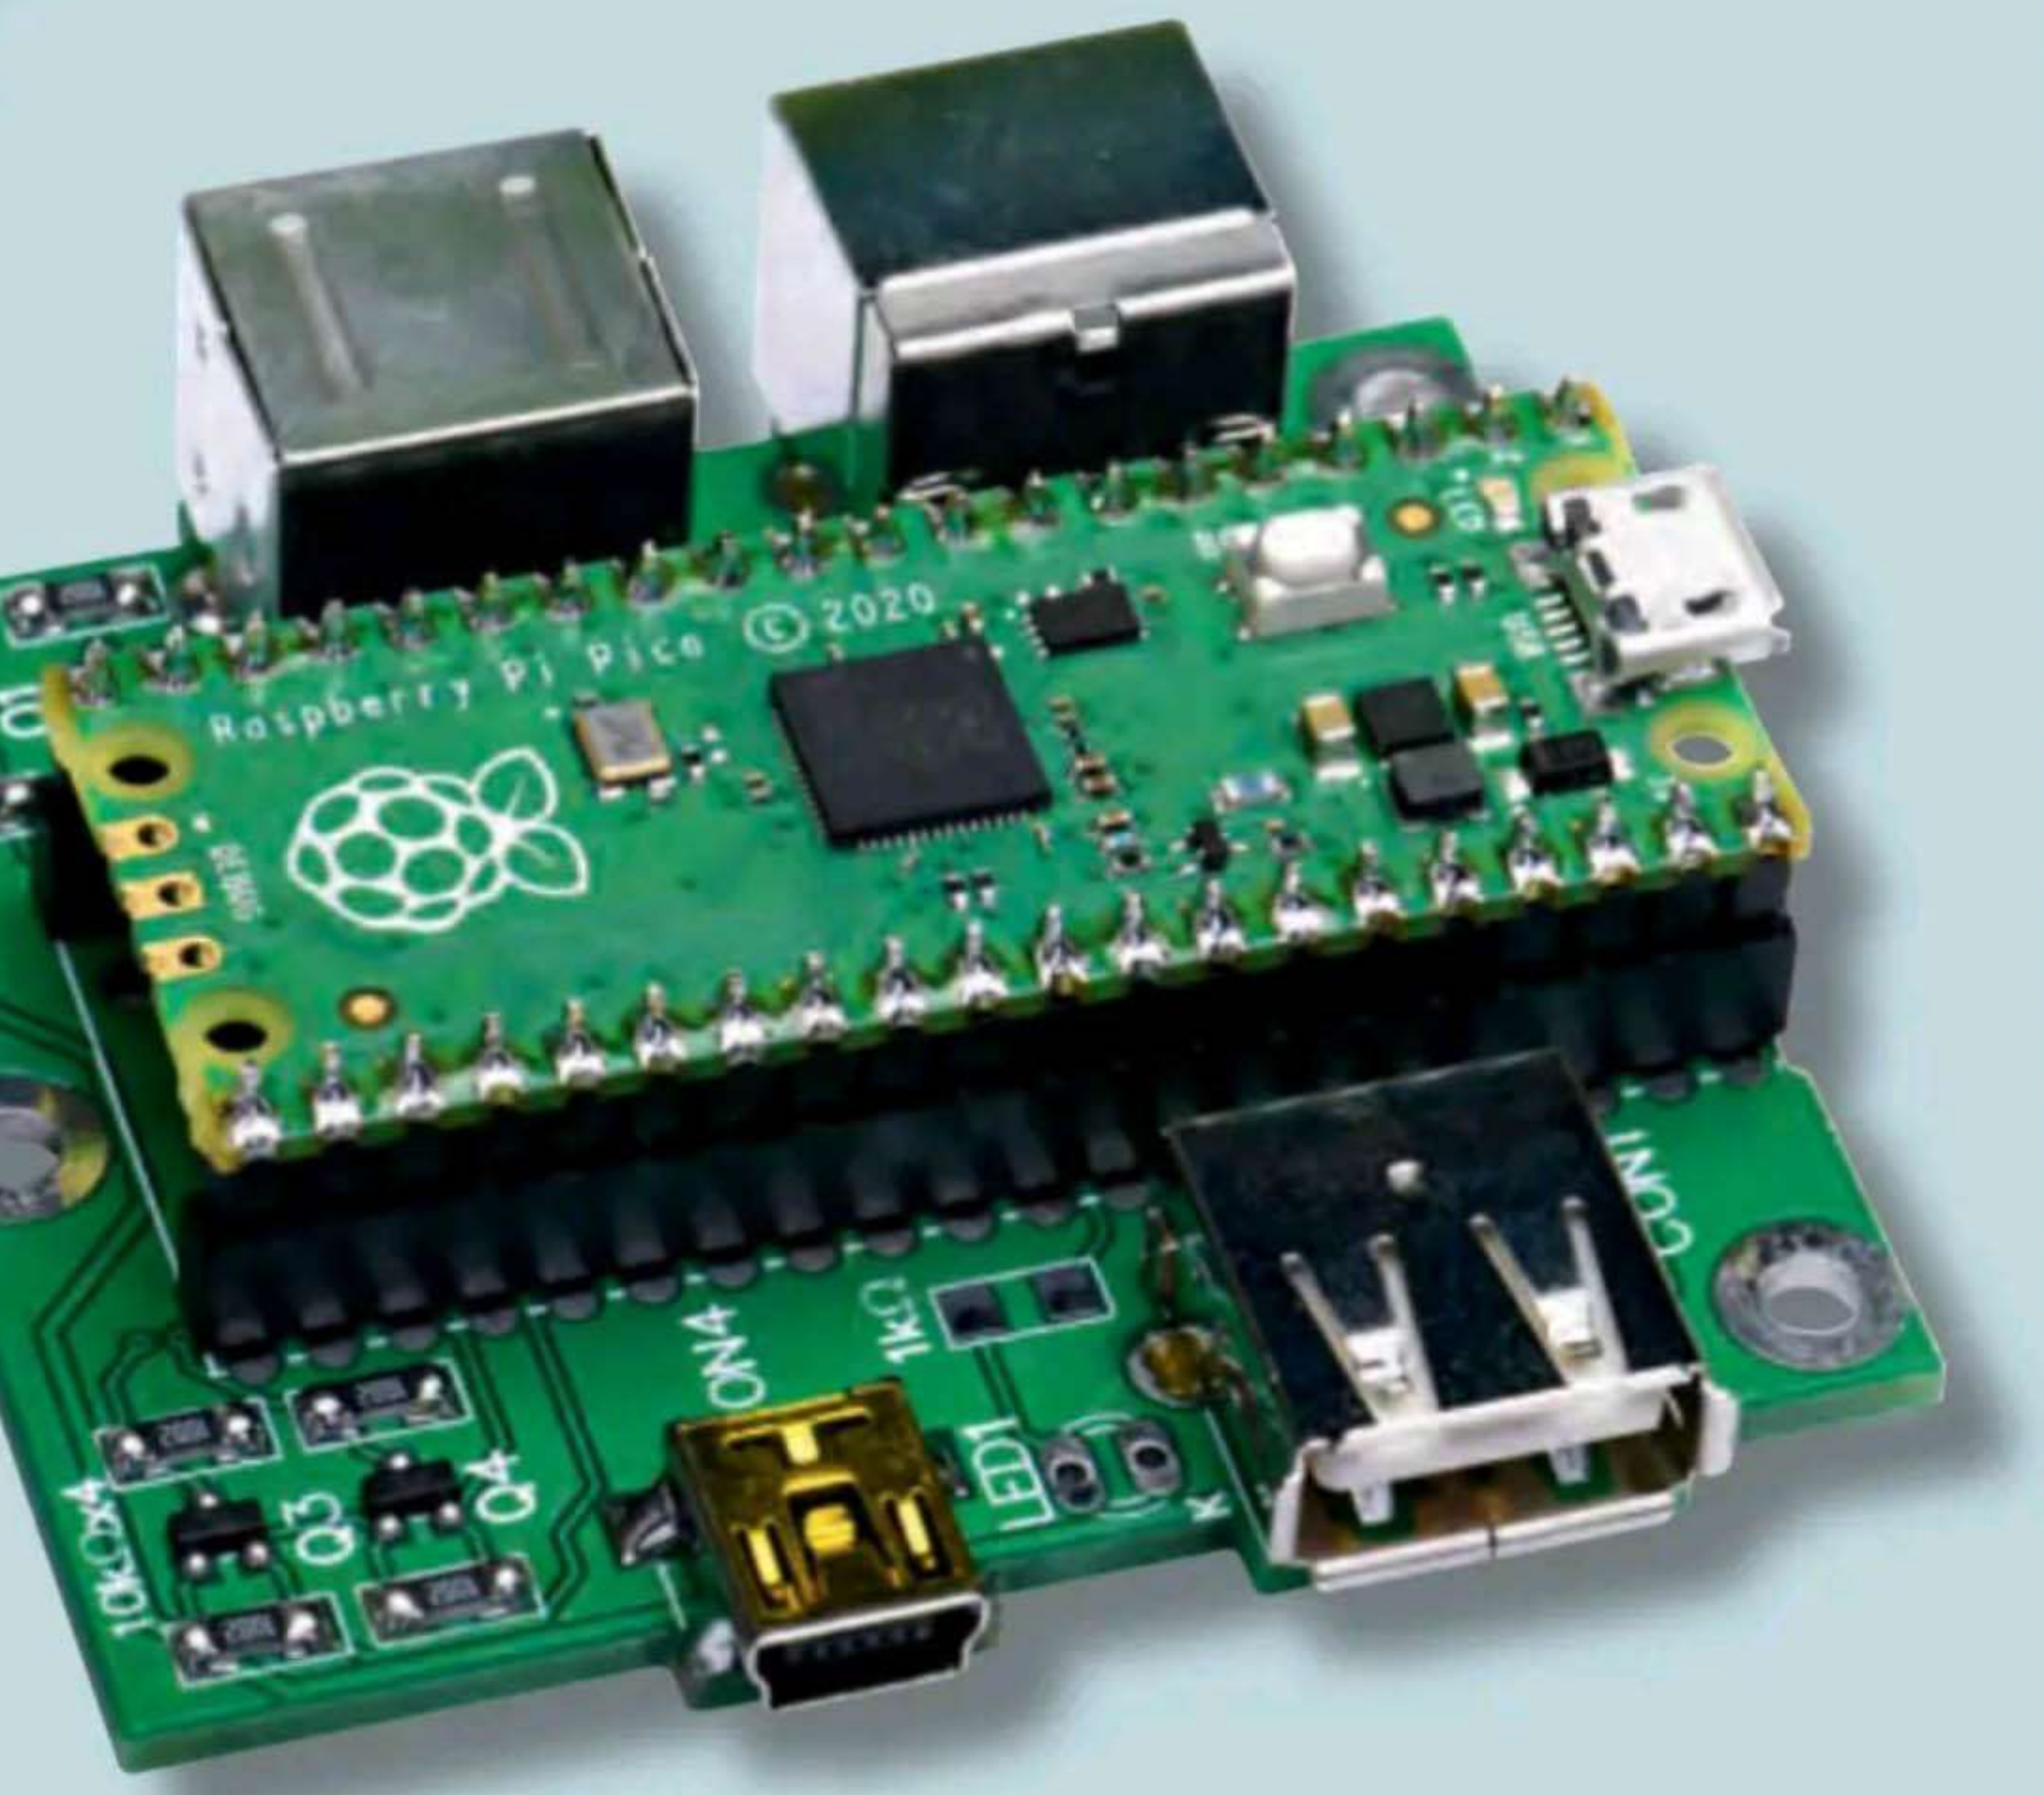

Convert USB to PS/2

Use a modern mouse and/or keyboard with legacy devices that only have PS/2 ports.

PLUS! All your favourite regular columns from *Max's Cool Beans* and *Circuit Surgery* to *Audio Out*, *Techno Talk* and *Net Work*.

On sale 5th of December, 2024

Content may be subject to change

Practical Electronics

Subscriptions hotline

To subscribe or renew online, paying by credit/debit card or PayPal, please visit our shop at: www.electronpublishing.com

For those wishing to pay by direct debit, cheque or over the telephone, please contact 'Select Publisher Services' by calling 01202 087631, or send a cheque (payable to: 'Practical Electronics') with your details to:

Practical Electronics Subscriptions, PO Box 6337, Bournemouth BH1 9EH, United Kingdom

Remember, we print the date of the last issue of your current subscription in a box on the address sheet that comes with your copy.

Digital subscribers, please call 01202 880299 or visit: www.electronpublishing.com

Published on approximately the first Thursday of each month by Silicon Chip Publications, U1 234 Harbord Road Brookvale NSW 2100 Australia. Printed in England by The Manson Group, Reynolds House, 8 Porters Wood, Valley Road Industrial Estate, St Albans AL3 6PZ. Distributed by Seymour, 86 Newman St, London W1T 3EX. Subscriptions UK: £33.99 (6mo); £59.99 (12mo.); £114.99 (2yrs). EUROPE: airmail service, £38.99 (6mo); £69.99 (12mo); £129.99 (2yrs). REST OF THE WORLD: airmail service, £46.99 (6mo); £84.99 (12mo); £164.99 (2yrs). Make payments to 'Practical Electronics', Practical Electronics Subscriptions, PO Box 6337, Bournemouth BH1 9EH United Kingdom. Email: pesubs@selectps.com. PRACTICAL ELECTRONICS is sold subject to the conditions that it shall not, without the written consent of the Publishers first having been given, be lent, resold, hired out or otherwise disposed of by way of Trade at more than the recommended selling price on the cover, and that it shall not be lent, resold, hired out or otherwise disposed of in a mutilated condition or in any unauthorised cover by way of Trade or affixed to or as part of any publication or advertising, literary or pictorial matter.

Did you know our online shop now sells the current issue of *PE* for £5.99 inc. p&p?

Practical Electronics

The UK's premier electronics, computing and maker magazine

Circuit Surgery
Processing signals digitally

All About Capacitors
Everything you need to know to pick the right capacitor for the job

60W DC load review
A low-cost device for testing supplies & batteries

Raspberry Pi Clock Radio & Media Player

Secure Remote Switch with up to sixteen keyfob remotes

The VGA PicoMite: A simple but capable BASIC computer!

PLUS!
Techno Talk – Tomorrow's World, today
Cool Beans – Driving LEDs with serial-to-parallel latches
Net Work – F. J. Camm, the practical man

Dec 2024 £5.99

www.electronpublishing.com | @practicalelec | practicalelectronics

Practical Electronics

computing and maker magazine

Practically Speaking
Avoiding electric shocks during repairs

up to 20 audio
reen or remote!

Nov 2024 £5.99

www.electronpublishing.com | @practicalelec | practicalelectronics

You read that right! We now sell the current issue of your favourite electronics magazine for exactly the same price as in the High Street, but we deliver it straight to your door – and for UK addresses we pay the postage. No need to journey into town to queue outside the newsagent. Just go to our website, set up an account in 30 seconds, order your magazine and we'll do the rest.

www.electronpublishing.com

Enhance Space and Defense Systems With SAMD21RT MCU

Small Footprint Radiation-Tolerant Arm® Cortex®-M0+

Are you looking for a solution to enhance your space and defense projects? Look no further than the SAMD21RT as this microcontroller offers a range of benefits, tailored to meet your demands.

Capturing the power of a 32-bit processor core, it enables development in the most confined spaces. With a compact footprint measuring only 10 x 10 mm, the SAMD21RT MCU delivers exceptional performance. Our goal is to design cost-effective, small package and power efficient products that will revolutionize your devices.

Our SAMD21RT microcontroller offers various resources and tools to benefit you:

- High level of radiation-tolerant with TID capability up to 50 kRad and SEL immunity up to 78 MeV. cm^2/mg
- MPLAB X Integrated Development Environment – expandable, highly configurable software program
- MPLAB® PICkit™ 5 in-circuit debugger/programmer
- SAM D21 Curiosity Nano Evaluation Kit

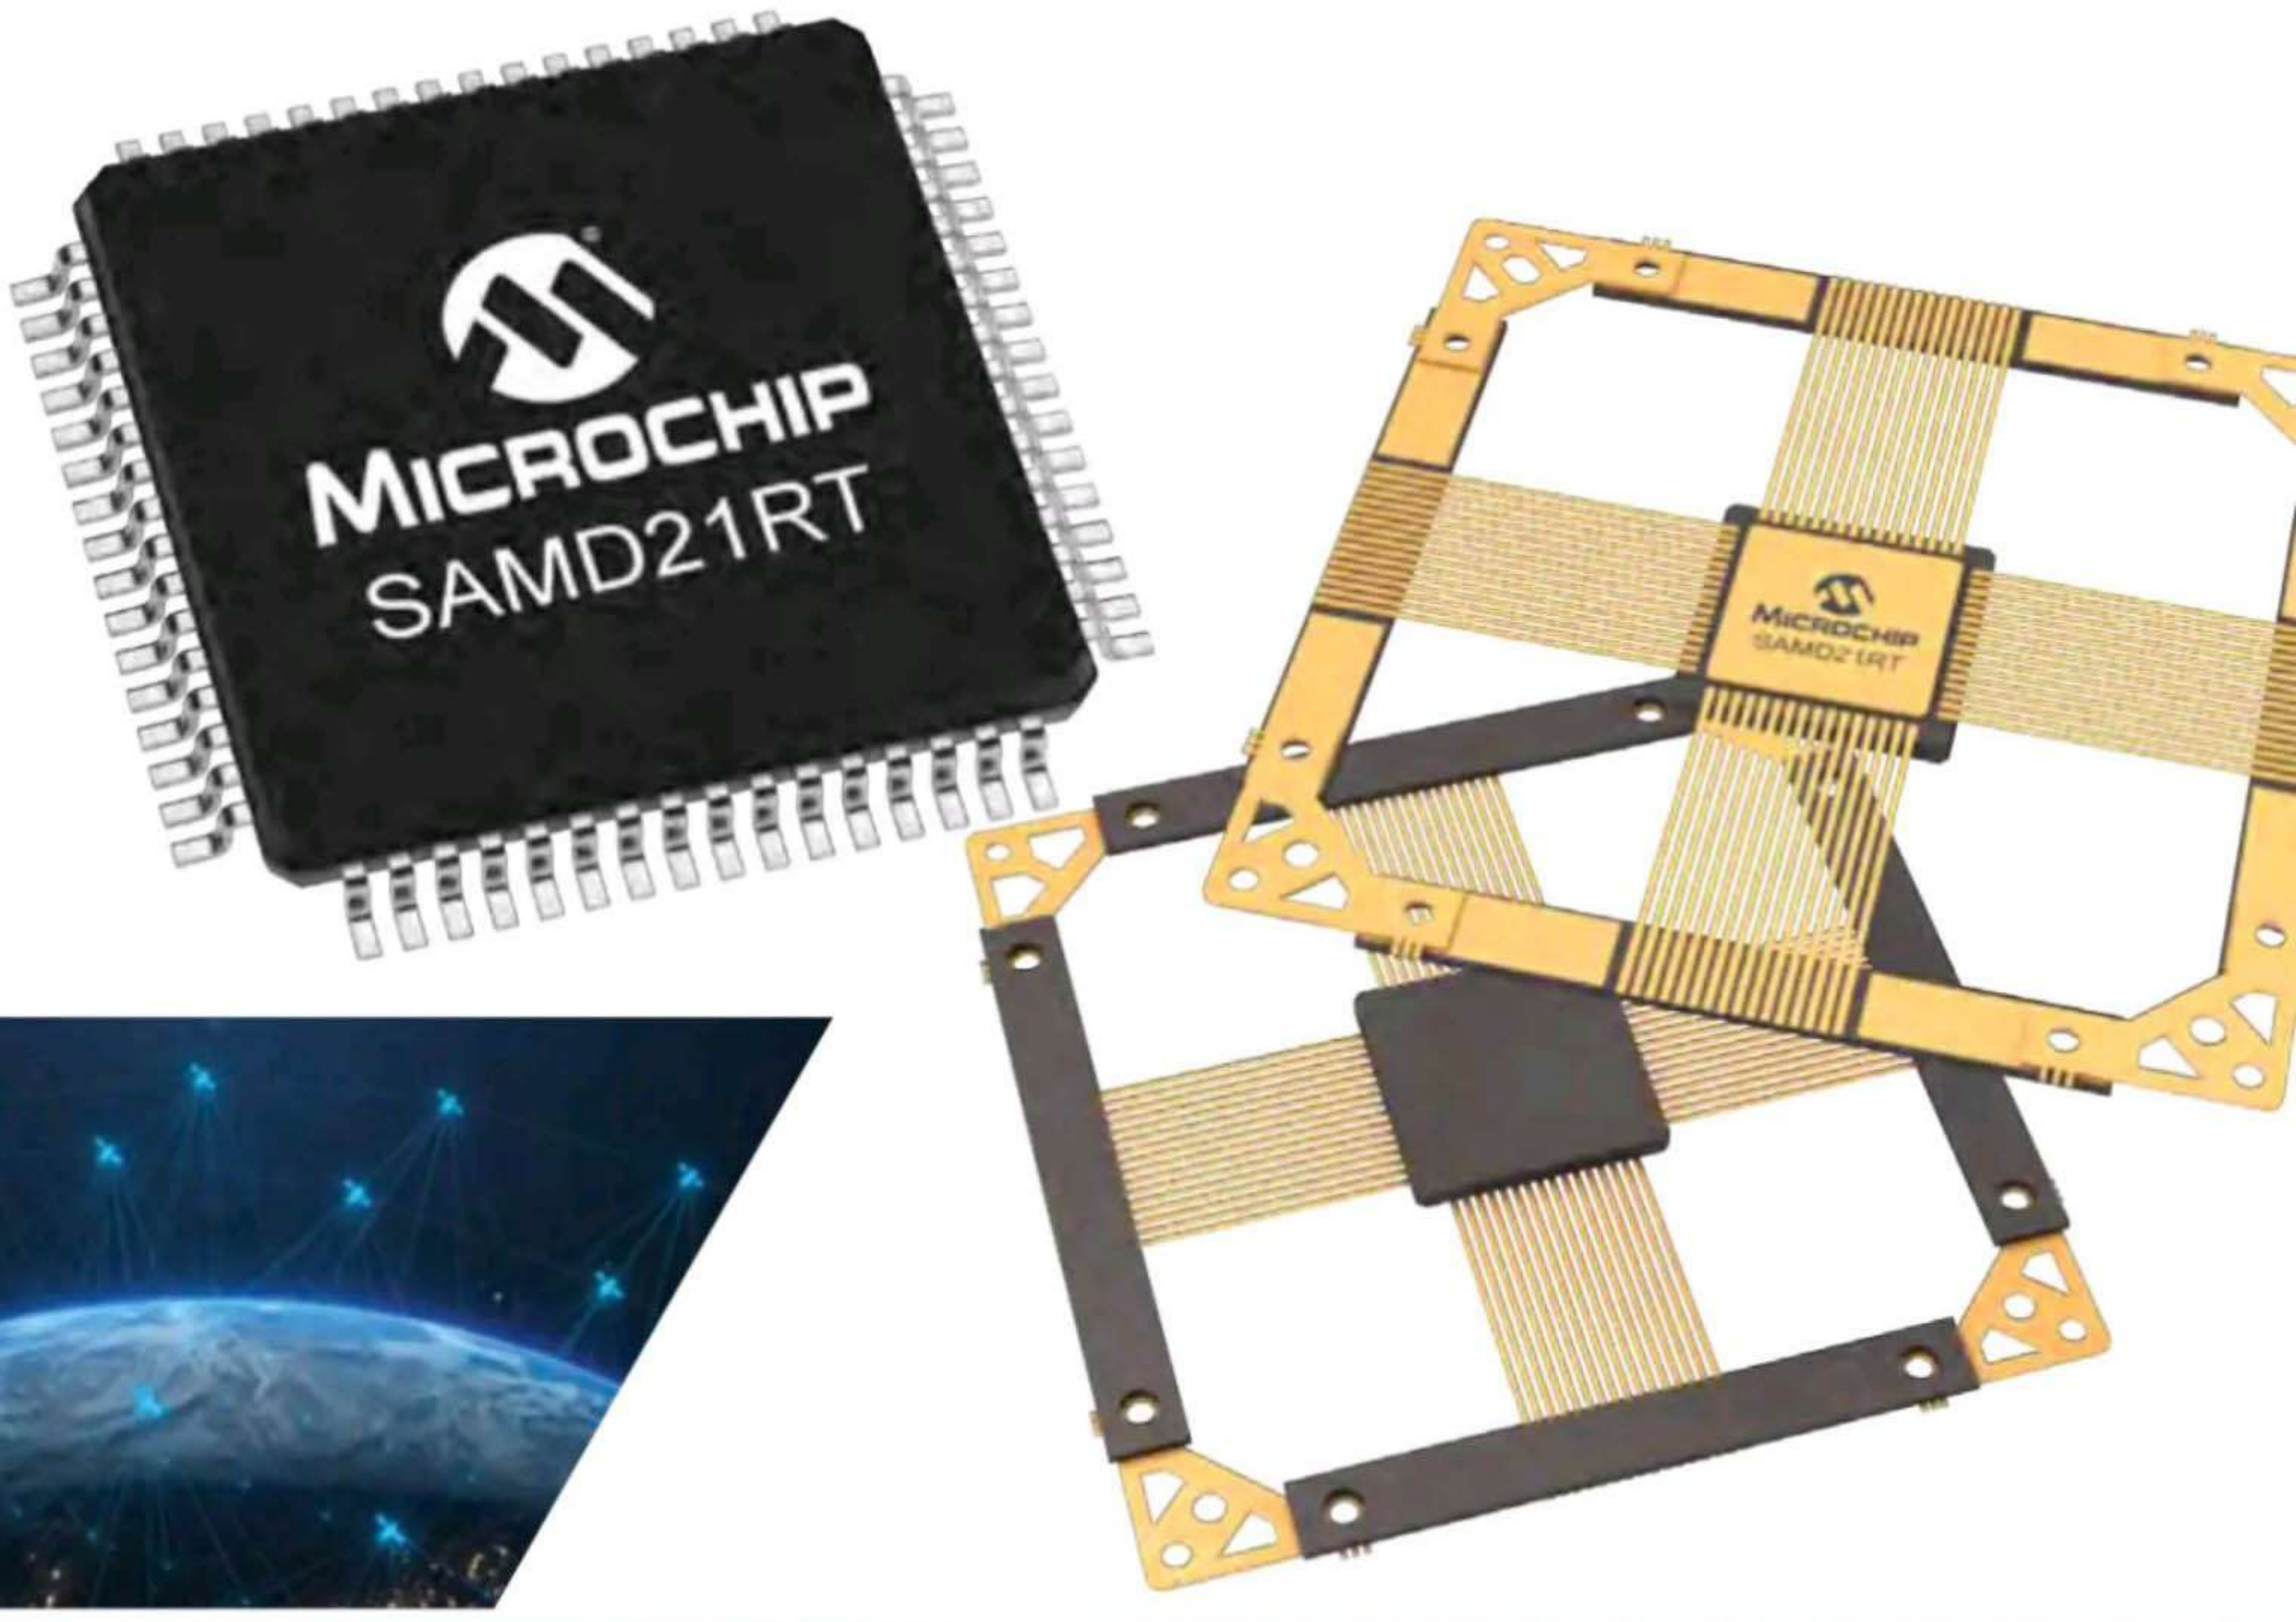

MICROCHIP

microchip.com/SAMD21RT

The Microchip name and logo and the Microchip logo are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries. All other trademarks are the property of their registered owners.
© 2024 Microchip Technology Inc. All rights reserved.
MEC2573A-UK-06-24